


Learning




A. CHARLES CATANIA

A. CHARLES CATANIA

University of Maryland Baltimore County

LEARNING



Catania, A Charles
Learning.

(Century psychology series)

Bibliography: p.

Includes index.

1. Learning, Psychology of. I. Title.

BF318.C37 153.1'5 78-21180

ISBN 0-13-527432-X

S.C.E R T., West Bengal

Date 22.8.84.....

Acc. No. 3045.....

THE CENTURY PSYCHOLOGY SERIES

154.4
CAT

©1979 by Prentice-Hall, Inc., Englewood Cliffs, N.J. 07632

All rights reserved. No part of this book
may be reproduced in any form or by any means
without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Prentice-Hall International, Inc., *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall of Canada, Ltd., *Toronto*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Prentice-Hall of Southeast Asia Pte. Ltd., *Singapore*
Whitehall Books Limited, *Wellington, New Zealand*

To Connie
To Bill
To Kenny
And even to Michael

CONTENTS

PREFACE ix

One

INTRODUCTION 3

The Language of Learning	4
The World and the Laboratory	7
Behavioral and Cognitive Psychologies	10
Implications of the Analysis of Behavior	13
Overview	14

Two

BEHAVIOR 17

Behavioral Hierarchies	18
Stimuli and Responses	21
Operations and Processes	24
Recapitulation	35

Three

ELICITATION 39

The Reflex	39
Properties of Elicited Behavior	41
The Temporal Structure of Elicited Behavior	48
Summary and Extension	59

Four**CONSEQUENCES OF RESPONDING: REINFORCEMENT 63**

The Law of Effect	64
Reinforcement	72
The Acquisition of Behavior	85
Summary	91

Five**CONSEQUENCES OF RESPONDING: AVERSIVE CONTROL 93**

Punishment	93
Reinforcement by Stimulus Removal (Negative Reinforcement)	104
Summary and Extension: The Language of Aversive Control	114

Six**OPERANTS: THE STRUCTURE OF BEHAVIOR 117**

Shaping: The Differential Reinforcement of Successive Approximations	118
Operants: Differentiation and Induction	119
Summary and Extension: The Structure of Behavior	136

Seven**DISCRIMINATED OPERANTS: STIMULUS CONTROL 139**

The Nature of Discriminated Operants	141
Discriminable Properties of the Environment	155
Summary	165

Eight**REINFORCEMENT SCHEDULES 167**

Variable-Interval and Variable-Ratio Schedules	169
Fixed-Interval and Fixed-Ratio Schedules	180
The Vocabulary of Reinforcement Schedules	183
Schedule Combinations	185
Summary: Schedule Analysis	196

Nine

RESPONDENT BEHAVIOR: CONDITIONING 199

- Conditional Reflexes 200
- Operant-Respondent Interactions 215
- Recapitulation: A Behavior Paradigm 221

Ten

VERBAL BEHAVIOR 229

- Functional Properties of Verbal Behavior 231

Eleven

PSYCHOLINGUISTICS 255

- Syntax: The Grammatical Structure of Language 256
- Semantics: The Meaning of Verbal Units 268
- Overview 273

Twelve

VERBAL LEARNING AND TRANSFER 277

- Verbal Learning Procedures 278
- Transfer 298

Thirteen

MEMORY 307

- The Nature of Memory 309
- Mnemonics 310
- The Varieties of Memory 313
- The Metaphor of Storage and Retrieval 325
- The Structure of Memory 334
- Summary 335

Fourteen**COGNITION AND PROBLEM-SOLVING 337**

Cognitive Processes 338

Problem-Solving 348

Fifteen**CONCLUSION 353**

Structure and Function 356

Learning and Evolution 358

REFERENCES 362**ACKNOWLEDGMENTS 395****INDEX 397**

PREFACE

Learning is central to the problems of psychology. To ask how much of what an organism does depends on its evolutionary history and how much depends on what it has experienced during its own lifetime is to ask what it can learn. Research on learning has ranged from relatively simple procedures with animals to the complexities of human language and problem-solving. Procedures in these areas are so different, and the literature so extensive, that the temptation is great to restrict one's attention solely to issues of animal behavior and learning, or solely to issues of human learning and memory. Many texts in learning yield to this temptation. In turn, the study of learning has become more and more divided, with each approach developing its own languages and research methodologies. In this book, part of my purpose has been to bring these literatures together, and to explore some of the continuities between human learning and the learning of other organisms. Humans are undoubtedly unique, but they share an evolutionary heritage with other species. The properties of animal learning are therefore likely to be relevant to learning in humans. Even if we only show that some kinds of human learning are not reducible to types of learning known to occur with other organisms, we have at least begun to define what is peculiarly human.

This book surveys the major areas in the psychology of learning from a consistent behavioral point of view. I will not attempt to outline the nature of a behavioral orientation here. That view has evolved considerably from its parochial origins, and is better treated in the context of specific psychological issues, some of which will be taken up later in the text. I will only note that to take a behavioral position one need not exclude from consideration aspects of human behavior sometimes called mental, such as thinking and imagining. Thus, topics often regarded as the exclusive province of contemporary cognitive psychology will be treated along with those more traditionally regarded as behavioral.

In its overall structure, this book has two major sections. Chapters 2 through 9 cover basic topics in animal behavior and learning, and chapters 10 through 14 cover those in human learning and memory. In some places, these sections are fairly independent, but more often the concepts developed in the first section are

prerequisites for the treatment of more complex issues in the second section. I have made a special effort to include examples of human behavior in discussing the relevance of animal learning studies, and to refer to appropriate concepts from animal behavior in discussing human learning and memory. The choices were to some extent dictated by the logic of the subject matter and by the availability of appropriate cases in the relevant research literature.

To study learning, one should know what learning is. In chapter 1, therefore, this book begins with the problem of defining learning. In that context, this introductory chapter also treats some aspects of technical vocabulary and of experimental method. Chapter 2 is organizational, providing an outline of the topics examined in detail in chapters 3 through 9. Chapter 3, on elicitation, covers the reflex and other relations generated when stimuli are presented to organisms. Other texts typically introduce conditioning at this point, but within the present organization that topic is more effectively deferred until later.

Chapters 4 and 5 show how the consequences of responding can affect behavior, in reinforcement and in aversive control. These topics raise the issue of classes of responses and classes of stimuli as behavioral units. They therefore lead to the concept of the operant, in chapter 6 on the differentiation of responses, and the concept of the discriminated operant, in chapter 7 on stimulus control. Some implications of these concepts are illustrated in chapter 8, which examines the ways in which behavior can be maintained over extended periods of time by reinforcement schedules. Chapter 9 then takes up the properties of conditioning and shows how they may be related to processes discussed in the preceding chapters.

In its behavioral treatment of language, chapter 10 provides a bridge to the complexities of human learning and memory. Some features of behavioral and cognitive approaches are explicitly compared in chapter 11, on psycholinguistics. These two chapters set the stage for an examination of verbal learning and transfer, in chapter 12, and memory, in chapter 13. Cognition and problem-solving, in chapter 14, present an opportunity for synthesis, because they bring together topics considered separately at various points throughout the text. The final chapter attempts an overview and integration of some major issues in the psychology of learning.

Throughout its history, the psychology of learning has been concerned with theories. Particular learning theories were developed, elaborated, and then displaced by others. Many remain with us, typically more restricted in their scope than when they were first introduced. Whatever the status of these theories, the empirical findings that gave rise to them are still to be dealt with. For that reason, this text emphasizes research findings rather than learning theories. The extent to which it is theoretical is mainly in its adherence to a consistent behavioral language and in its concern with a systematic organization that can accommodate the various procedures and processes of learning. Although theory is not emphasized, I have tried to include sufficient information about experimental procedures, terminology, and data to provide an effective starting point for the student, instructor, or general reader who wishes to pursue particular theories. The emphasis of the book is not so

much on the interpretation of particular findings as on the relations among the varied phenomena included within the psychology of learning. My intent has been to make the book useful not only to those who are already behaviorally inclined, but also to those who are decidedly not of the behavioral persuasion.

This book has been an effort of some years. Its content grew over successive offerings of an introductory lecture course in the Psychology of Learning, first at the University College of Arts and Science of New York University and then at the University of Maryland Baltimore County. I am indebted to my students and colleagues at both campuses; I must especially mention Robert E. Silverman, T. James Matthews, Barry Schwartz, Kenneth Keller, Rona Yohalem, Jacob Gewirtz, and Eliot H. Shimoff. As my teachers and colleagues, numerous others contributed by their comments, discussions, and encouragement. I name only the following, mainly because I can still identify particular contributions of each: Leonard Cook, Willard F. Day, Lewis R. Gollub, Eliot Hearst, Ralph F. Hefferline, Richard J. Herrnstein, Herbert M. Jenkins, Victor G. Laties, J. A. Nevin, Howard Rachlin, George S. Reynolds, B. F. Skinner, William C. Stebbins, S. S. Stevens, Stephanie B. Stolz, Herbert S. Terrace, and Thom Verhave. In such a listing, omissions are inevitable, but to a singular group of forty-odd unnamed others it will come as no revelation that at least six names remain deliberately secret. I wish also to express my appreciation for the invaluable help of Geraldine Hansen, Jill Himmer, and Lee Ratner, in preparation of the manuscript, and of Jack Burton, Neale Sweet, Ted Jursek, John Isley, Alison Gnerre, and Paul Spencer, in guiding the book through its successive stages from conception to publication.

Above all, however, one further acknowledgment remains. I cannot say how much I owe to the teaching and the writing of Fred S. Keller and W. N. Schoenfeld. Their courses and their *Principles of Psychology* introduced me to the analysis of behavior. I hope that this book is true enough to their teaching that they will recognize something of themselves in its pages.

A. Charles Catania
Columbia, Maryland
October 1978

Introduction

THE LANGUAGE OF LEARNING

THE WORLD AND THE LABORATORY

BEHAVIORAL AND COGNITIVE PSYCHOLOGIES

IMPLICATIONS OF THE ANALYSIS OF BEHAVIOR

OVERVIEW

The English word *learning* comes from an Indo-European root, *leis-*, which meant a track or furrow. Before it reached its present form, it went through many changes: *læstan*, *leornian*, *lernen*. At various times in the evolution of our language it might have been understood as following a track, continuing, coming to know, or perhaps even getting into a rut. From the same root came the verb *last*, to endure.

The study of learning is concerned with how organisms come to behave in new ways. But behavior can change in many ways, and some ways may be more appropriate as instances of learning than others. For example, two children each may have learned to answer multiplication questions. If one learned by studying the multiplication table and the other by mastering an electronic calculator, we would certainly agree that the children each learned something. We would also say that what each learned was different, and we might argue further that only the first child really learned to multiply. After all, that child could answer multiplication questions in a variety of circumstances, whereas the other child could do so only with a calculator. Thus, we could say what each child learned by describing what the child now can do and the circumstances in which the child can do it.

This does not define learning, however. The task of definition is complicated, but the word *learning* does not give us trouble in everyday conversation. We can usually say whether or not we have learned something, and we are usually able to agree on what counts as learning even if we cannot give a definition. We may therefore study learning by identifying particular cases, by observing the conditions under which they occur, and by examining how they can be combined. In our study of learning, we must remember to distinguish between two different problems: first, what is the nature of the phenomena or events that we speak of as learning, and second, what is the best way of talking about them?

THE LANGUAGE OF LEARNING

First consider our language. We regard words like *learning* and *knowledge* as important and obvious. But often we fail to distinguish among the different functions of these words, and this can cause confusion. Sometimes we speak of learning about something; at other times, we speak of learning how to do something. Someone who has learned how a bicycle works will not necessarily know how to ride a bicycle; inversely, someone who has learned to ride a bicycle may be unable to say what keeps a bicycle up, or even how the gears or the brakes of a bicycle work. We must decide whether we will treat these two kinds of learning together or separately. Philosophers are concerned with this kind of distinction when they debate the difference between "knowing how" and "knowing that" (e.g., Ryle, 1949).

If *learning* could be defined in a sentence or two, we would have no problem. We would define the word and then discuss the conditions under which learning takes place, the kinds of things that are learned, the limitations of learning, and so on. But learning has come to mean different things at different times and to different people, and it is not so easily defined. In fact, we might regard parts of this book as explorations of appropriate definitions for *learning*. What should we count as learning, and what should we exclude?

Consider some examples. A pigeon discovers food in its travels, and returns to that place on later occasions when it is hungry. A child becomes able to read a story or to spell simple words. A dog is taught to sit or lie down on command. A patient who once had a bad experience in a dentist's office consistently feels afraid in the waiting room. A young cat, after its early hunting expeditions, comes to avoid skunks and porcupines. A student reads a chapter in a mathematics textbook, and finds a way to solve a problem that once was baffling. A voter reads about the candidates for office in a newspaper, and votes on the basis of that reading several days later.

What do these examples have in common? Whether they involve dogs and rats or children and adults, each is an example of a change in what an organism does in particular circumstances. Learning, in fact, is often defined as a relatively permanent change in behavior resulting from experience (e.g., cf. Kimble, 1961, pp. 1-13; K. Smith, 1969). But what do we mean by "relatively permanent," by "behavior," by "experience"? Staring at an eclipse of the sun is undoubtedly an experience, and it will certainly alter one's future behavior if it permanently damages the eyes. But if someone claimed that this was an instance of learning, we would probably disagree. The preceding definition, then, is unsatisfactory. Yet we would probably agree that each of the earlier examples was an instance of learning. We might wonder, however, whether there is any profit to classing a pigeon who learns a route to food together with a student who discovers a solution to a mathematics problem.

Perhaps the issue can be resolved by turning to an organism simpler than the pigeon or the human. Consider the single-celled slipper-shaped micro-organism, *paramecium*. Whether single-celled organisms can learn is one of the lesser but

persistent controversies of the psychology of learning. In one experiment, a food-coated needle was repeatedly dipped into the center of a water dish that contained paramecia. After each feeding, the water was stirred, but nevertheless the paramecia began to congregate, after several feedings, in the region where food was presented. It was therefore argued that the paramecia had learned the location of food. The argument ended, however, when it was found that the stirrings did not adequately disperse the physical residues of past feedings. The congregation of the paramecia depended on these residues rather than on past opportunities to eat (Jensen, 1957). The experiment, therefore, had not demonstrated learning.

In another experiment, paramecia were allowed to swim into a lighted area that had also been made hot. After this pairing of light with heat, the paramecia avoided other lighted areas even though those areas had not been heated. It was argued that, through the pairing, avoidance of heat had been conditioned to the light. But this argument, too, abated, when it was later discovered that heated paramecia avoid light whether or not the light and the heat have been paired with each other (Best, 1954). The paramecia avoided the light because that is what heated paramecia do, and not because they had learned how light was related to heat in these experiments. This experiment, too, had failed to demonstrate learning. Thus, the answer to the question, whether paramecia learn, has changed from time to time.

But what kind of a question is this? Certainly the behavior of a paramecium can change. Yet the same change in behavior might or might not be considered an instance of learning, depending on the circumstances that led up to it. The question is in part about paramecia, but it is also about the conditions under which we say that an organism has learned.

Consider another example. If a paramecium is exposed for a while to a temperature somewhat higher than the one at which it had been living, it may subsequently survive still higher temperatures that would otherwise have killed it (Beale, 1953). Might we not say that the paramecium had *learned* to tolerate higher temperatures? And what could we discover about the way in which this temperature adaptation works that would convince us to speak differently about it? There is no simple answer to these questions, but they at least illustrate that we must worry not only about the behavior that we study but also about the way in which we speak of it.

At this point, some might argue that these paramecium examples are irrelevant in any case. The paramecium is a single-celled organism, and therefore it can have no nervous system. It might be suggested that our problems of definition would be resolved if we added that the change in behavior had to come about through some kind of change in the brain. But we do not look at an organism's brain to decide whether it has learned something. We all have learned to say when we or others have learned something, but few if any of us have ever seen a brain doing anything. Even if we could watch a brain doing something, how would we know that what it was doing was learning? This is not to say that learning has no physiological basis; it would be interesting to know what physiological changes accompany learning. Yet if we cannot even say what learning is, we might find it

difficult to know how to study its physiology. In fact, we might expect an adequate physiology of learning to depend on an adequate understanding of the behavioral properties of learning.

The assertion that learning must take place in the brain carries with it many assumptions. If we study what an athlete learns, for example, shall we assume that all of the important changes are in the nervous system and none in the muscles?¹ What reason would we have for excluding muscle? And once we recognize the importance of muscle in athletic learning, why not also in other kinds of learning? Would we then be ready to look for learning in bones or blood as well?

Overstatement would be easy at this point. We could argue that the brain is irrelevant to behavior and defend this position by citing hypothetical cases. For example, it is likely that if we removed the tip of an attractive woman's nose, we would change her behavior far more drastically than if we removed an equal volume of tissue from the cortex of her brain. We could then say that, because the former operation would dramatically affect her social behavior whereas the latter might not even be detectable without special testing, social behavior resides in the nose and not in the brain. That would make the point, but the example is admittedly frivolous. We must not take it seriously. We can readily concede that changes in behavior are related in important ways to activities of the nervous system. At the same time, we must remember that it is the whole organism that behaves. Sensory systems keep the organism in contact with the world, and muscle systems allow the organism to act upon the world in turn. A nervous system that had nothing to do with either of these systems would be of little interest. For this reason, we will be mainly concerned with the behavioral properties of learning, and not with its physiological basis.

The problems so far have only partly been about the facts of learning; they have also been about our language. Languages are changeable, and their vocabularies reflect what is currently important to their speakers. One trouble is that the language that has evolved in our everyday dealing with others is not necessarily the most suitable for a language of learning (that is one reason for the etymologies or word histories included at the beginning of each chapter). We are usually more interested in what other people know and in what they are likely to do than in how they came to be that way. A parent or a teacher, for example, might become concerned if a child consistently fought with other children instead of playing cooperatively. But once the child is taught to play cooperatively, the adult might not care whether the child came to be that way because cooperation was praised or because anything else that the child did was punished.

Our language for describing what people do is useful. It is important to know what to expect of others. Thus, we describe people in terms of how they are likely to behave. We speak of each other as outgoing or reserved, easy-going or compulsive, trustworthy or unreliable. Describing people with words like artistic, athletic, social, intellectual or musical specifies their preferred activities. This kind of vocabulary, however, is not suitable for discussing how particular interests or traits

¹It is interesting that human muscle cells, like human nerve cells, neither multiply nor regenerate in maturity; with exercise, individual muscle cells grow in volume (Edgerton, 1973).

developed in an individual. For example, there is an important difference between lying and telling the truth. But the habitual liar may lie for some of the same reasons that an honest person tells the truth. If circumstances are such that one child keeps out of trouble only by telling lies and another only by telling the truth, it would not surprise us if the first child grew up to be less truthful than the second. Yet the behavior of each child is shaped by its consequences; each child behaves so as to keep out of trouble. This shaping of behavior is our interest, but our everyday vocabulary does not equip us well for discussing it.

Similar problems exist in fields other than psychology. When physicists look at events in the world, they do not find the everyday vocabulary adequate. They coin new terms or take over existing ones. The latter course can create difficulties. Words like *work*, *force*, and *energy*, for example, mean different things to physicists in their technical talk than they do to most other people in casual conversation. Fortunately for physicists, many of the phenomena that they study are far enough removed from our daily experience that we do not confuse their technical language with our ordinary discourse.

This is not the case in psychology. We are all inescapably concerned with behavior. We speak of how people grow and change, we speculate about the reasons people have for doing things, and we ourselves learn new facts and acquire new skills. If we wish new ways of talking about these events, we must take care that the new language does not get confused with the old one. We all have spent most of our lives talking in particular ways about what we do, and these familiar ways of talking may interfere with any new ways of talking that we may try to establish. Some parts of this volume will be concerned with establishing a behavioral language, and it will be important to note that the language will not be simply a paraphrase of everyday usage; it will instead require some new ways of dealing with familiar phenomena or events.

THE WORLD AND THE LABORATORY

The problems of language are made even more difficult because we live in a complex world. The events that influence what we do and how we change do not occur in isolation. Thus, to understand any given situation, we must analyze it; we must strip away the unessential details to get at the basic processes that may be common to different circumstances. To analyze something is simply to break it down into its component parts. For this reason, we turn to the laboratory. We begin with the study of organisms simpler than ourselves, in simple environments. We must face the objection, of course, that a laboratory learning experiment is artificial and therefore inappropriate for establishing generalizations about learning outside of the laboratory. But we must start with simplicity so that we can develop techniques and vocabularies with which to deal with complexity.

The controlled laboratory environment enables us to look at one thing at a time. We can arrange circumstances so that we know what goes into the experimental situation; if we are careful, we can exclude the distractions that might otherwise obscure the processes that we wish to study. The simplicity of our

laboratory environment may also help us to see important aspects of learning, and therefore to develop a vocabulary appropriate to our subject matter. We have to be able to identify events before we can study their properties.

These advantages of the laboratory are illustrated in the following example about a pigeon that earns its food by pecking (cf. Kelleher and Gollub, 1962). The account is of interest because the outcome is not at all intuitively obvious. In itself, it is not a learning experiment, but it grew out of experiments on how organisms learn to order their behavior in time, and its methodologies have proved useful in other learning studies. Imagine a pigeon in an experimental chamber. The pigeon has been trained to earn food by pecking a key, a translucent disk recessed in the chamber wall (in recording pecks, this key operates much like a telegraph key). The key can be lit from behind to present visual stimuli to the pigeon. Currently, it is lit white. Our equipment is set so that every two-hundredth peck gives the pigeon access to food in a hopper below the key. We reached this point gradually, first by letting every peck produce food, and then only every fifth peck, and then only every tenth, and so on, until finally the pigeon was pecking 200 times for each food delivery.

In this situation, with no food but what is produced by pecking, the pigeon easily maintains an adequate daily ration. In fact, the pigeon eats its daily diet in as little as 4 or 5 minutes. If each food delivery lasts only 5 seconds (including the time to get from key to food hopper), the pigeon may be finished eating for the day after only 60 food deliveries. Even so, this pigeon must peck at least twelve thousand times daily to earn this much food.

Technically, the requirement of 200 pecks per food delivery is called a fixed-ratio schedule of reinforcement: *fixed* because the required number is constant from one food delivery to the next, and *ratio* because the schedule specifies the ratio of pecks to food deliveries (in this case, 200 to 1). The pecks occur because they produce food; we therefore say they are *reinforced* by food. But we do not have to reinforce every peck to keep the pigeon pecking, and the schedule simply specifies how we choose the particular peck that will produce food. Reinforcement schedules will be considered in detail in chapter 8, however, so let us return to the pigeon.

Food has just been presented, the hopper has moved out of the pigeon's reach, and the pigeon is standing before the white-lit key. Time passes, perhaps only 15 or 20 seconds or perhaps several minutes. Eventually the pigeon approaches the key and begins to peck. Within a few pecks, the pigeon is pecking as rapidly as 5 to 10 pecks per second, and is not likely to pause at all as it completes its 200 pecks. The last of these pecks operates the food hopper. Given its rapid pecking, the pigeon may peck 2 or 3 times more before it stops and moves down to the hopper to eat. When the food is withdrawn, time again passes before the pigeon returns to the key and begins its next 200 pecks. This is typical fixed-ratio performance: a pause after reinforcement, followed by rapid and uninterrupted pecking that ends with the next reinforcement.

The point is that this is easy to get the pigeon to do. Our pigeon will earn enough food daily to keep itself alive and healthy for an indefinite time (its life expectancy is at least 15 years). If we kept it in its chamber 24 hours a day, we

would not have to intervene except to provide water, clean the chamber and keep the food hopper filled.

But pigeons cannot count very well, and we might wonder whether the pigeon is at a disadvantage each time it goes through its 200 pecks. Suppose we decide to help it to estimate how many pecks are left in the sequence, by presenting different lights on the key depending on the number of pecks that have occurred. Let us arrange that, after each food delivery, the key is blue until the pigeon has pecked 50 times. The fiftieth peck turns the key green. After 50 more pecks, the key turns yellow. Then, after 50 more, the key turns red. After 50 more pecks with the key now red, the 200 pecks are completed and the last peck operates the food hopper. The pigeon eats, and the key is once again blue. The number of pecks required for food remains the same; only the lights on the key have changed. (The pigeon, by the way, can easily distinguish among these colors.)

With no distinctive stimuli, the pigeon worked well enough, earning an adequate daily ration. Have we helped the pigeon by providing the colors to tell it where it is in the sequence? The surprising effect of adding these stimuli is that the pigeon begins to slow down. The pauses after each food delivery, when the key is blue, become longer. When the pigeon finally begins to respond, the pecks, which used to occur in rapid succession, now occur sporadically. When the key turns green, after the fiftieth peck, the pigeon may even pause before starting the next 50-peck sequence. When the key had been always white, the pigeon did not pause after 50 pecks; the 200-peck sequence usually occurred without interruption. In green, another 50 pecks, and the key turns yellow. This time the pigeon is less likely to pause at all: 50 more pecks turn the key red, and the pigeon quickly completes the last 50 pecks, food is delivered and the pigeon eats. But then the key is blue again, and another long pause begins.

The added colors did not help. The pigeon takes much longer to earn each food delivery than it had when the lamp remained white. In fact, even though this pigeon had maintained an adequate diet at 200 pecks per food delivery when the key was always white, we would now be wise to watch it carefully to be sure that its daily food intake does not decrease substantially. And this results not from any change in the pecks required for food, but simply from a change in stimuli. When we added the colors, we broke down the 200 pecks into 4 distinct units of 50 pecks each. We call these *chained* fixed-ratio schedules: the separate stimuli correspond to the links of the chain. But the chain does not hold the 200-peck sequence together more cohesively; instead, it breaks it up. When the key was always white, pecking at the start of the sequence was not so very different from pecking near the end, when a peck finally produced food. Once the stimuli were added, however, pecking in the presence of the early stimuli became less like pecking later. In blue, for example, pecks never produced food; at best they turned the key green, but pecks never produced food during this color either. With pecking reduced early in the chain, the time to complete each 200 pecks increased. In chained schedules, a stimulus supports less responding the further it is from the end of the sequence. Even severe deprivation of food is not sufficient to counteract this effect.

There are other ways to alter the pigeon's performance. For example, increas-

ing the required pecks, to 300 or even 500 per food delivery, also slows the pigeon down, creating long pauses after reinforcement and frequent interruptions of pecking; this effect is called *ratio strain*. We could now go on to examine the phenomenon in more detail, separating those properties of the situation that mattered from those that did not. For the present purposes, however, the point is that the outcome is not obvious. It arises out of an experimental analysis of behavior and, because this effect occurs with other organisms besides the pigeon, with different kinds of responses and different response requirements, and with different kinds and orderings of stimuli, it has substantial generality. Thus, we may wonder whether it is relevant to human behavior.

The things people do, of course, depend on consequences more complex than food deliveries. Any account of human behavior must grapple with the fact that completing any sequential task, such as writing a paper, preparing a meal, or finishing a research project, is affected not only by its properties as a behavior sequence, but also by competing obligations and interests, delays between actions and outcomes, and any number of other factors. Nevertheless, so much of what we do involves sequential behavior that we must wonder whether we might sometimes operate under the strain of too many links in our chains. The ability to formulate and achieve long-range objectives is supposed to be a unique characteristic of the human species. But in the pigeon's behavior, adding even a single link to a chain can have devastating effects. Perhaps we ought not assume that human behavior will inevitably be immune to such effects.

In any case, the example demonstrates that the laboratory may yield interesting phenomena even under relatively simple circumstances. The pigeon's chamber with its single key and its feeder may seem austere, and yet merely adding some colored lights produces effects with perhaps far-reaching implications. The controlled laboratory environment, therefore, reveals properties of behavior to us at the same time that it helps us to develop a language that is appropriate to this subject matter.

BEHAVIORAL AND COGNITIVE PSYCHOLOGIES

What is this subject matter? Sometimes we speak about what people do; sometimes we speak about what they know. Even psychologists do not agree on whether the subject matter of psychology is behavior or the mind. On the one hand, whenever we study learning we know that what an organism does is the only thing accessible to us. There is nothing else for us to examine but behavior. A human in a learning experiment may describe thoughts or feelings, but these descriptions are still only behavior; verbal behavior may be of a special sort, but it is behavior nonetheless. No matter what the phenomena we study in psychology, our terminology and observations and theories must ultimately be derived from behavior, from what organisms do.

On the other hand, there is more to an organism than shows in its behavior. Two students may sit silently through a lecture, and yet it may be clear to the instructor that one is able to answer certain questions and solve certain problems

while the other cannot. Although they might earlier have been distinguishable on the basis of past performances, the fact remains that they are not currently behaving differently. The difference is in what each potentially can do; we might say simply that one student knows more than the other. When we study this knowledge, it is tempting to say that we study the student's mind.

The debate between psychologists who call themselves behaviorists and psychologists who call themselves mentalists or cognitivists has been a long-standing one. To some extent, the debate has been about appropriate ways of talking of psychological events. The behaviorist argues that, because behavior is all that is available to measure, the language of mental events misleads by calling attention away from the fundamental subject matter. The behaviorist is especially troubled when a mentalistic account is accepted as explanatory and therefore discourages a further inquiry. For example, we may sometimes say that an idea, a feeling or a hunch led someone to do something. The behaviorist does not dispute the existence of ideas, feelings and hunches, but rather criticizes their invocation as causes of behavior. It is too easy to be satisfied with an explanation in these terms; it is not enough to say that someone did something because of an idea, a feeling or a hunch. Ideas, feelings and hunches are about the world, and therefore must have their origins in our experiences with the world. We must look further, to these past experiences or, in other words, to past behavior, to account for what we do. If we are successful, we will at the same time have something to say about the origins of our ideas, feelings and hunches. We will be able to say that these are names for certain things that happen to us when we deal with events in the world.

The cognivist maintains that such a view is unnecessarily narrow. There must be processes that occur in our dealings with the world that are not observable in our behavior. When we try to recall a word that is "on the tip of our tongue" or try to solve a problem by "sleeping on it," things are happening that do not show in our behavior, and we may not even be able to report these things in our verbal behavior. If we can find out something about such processes, it cannot help but be relevant to our study of learning.

Yet much of the dispute between these two kinds of psychologists is simply a matter of language. The difficulties have persisted not because behaviorists and cognitivists cannot understand each other, and not because there are psychological problems that either can resolve that the other cannot, but rather because the two kinds of psychologists are interested in different types of questions. The behaviorist tends to be interested in questions of function, and the cognivist in questions of structure (Catania, 1973c).

Suppose we are interested in teaching a child to read. On the one hand, we may become concerned with what we have to do to involve the child in reading. We worry about what will keep the child alert, what will make the child pay attention to the words presented, and what will help the child remember what the various words are. Will we be more successful if we reward the child for being correct or penalize the child for being wrong? When we arrange different consequences for different answers the child might give, we are determining the functions of these various answers or, more precisely, the functional relations

between behavior and its consequences. Problems of function are concerned with how behavior and environmental events can act upon each other.

On the other hand, no amount of concern with the effects of reward and punishment on the child's mastery of reading will tell us the most efficient way to present reading materials to the child. How is reading structured? What is the best way to order the materials? We might debate whether we should teach the child to read by starting with single letters, with syllables, or with whole words. When we arrange different orders of presentation of the materials to be learned, we are concerned with the effective structural relations within the subject matter. Are words unitary structures, or are they complex structures built up from simpler units, letters or syllables? Problems of structure are concerned with how behavior and the environment are organized.

Both problems are important. Any attempt to affect how children learn to read will be deficient if it ignores either of them. But it is also important to distinguish between them. There is little value in arguing over the priority of one or the other approach if both are indispensable to our understanding. Historically, difficulties arose because psychologists interested in functional problems tended to speak a behavioral language, whereas those interested in structural problems tended to speak a mental or cognitive language. But a behaviorist can study structural problems just as a mentalist can study functional ones. That these two different orientations, behaviorism and mentalism, tended to be correlated respectively with two different research emphases, functional and structural, was not inevitable. It is easy to see, however, how the correlation may have come about. If the experimental concern is functional, one studies the consequences of particular relations between specified environmental events and specified actions; these relations are conveniently expressed in the behavioral language of stimuli and responses. If the experimental concern is structural, one studies the properties of particular capacities or abilities; these properties are conveniently expressed in the cognitive language of knowledge and mind.²

We need not be side-tracked by this controversy in psychology. We will consider both functional and structural problems in learning and will therefore examine both types of research. This text is written from an explicitly behavioral orientation, but in dealing with problems of structure it will necessarily extend this orientation to cognitive concepts. Although these two psychological orientations differ in their languages and in the research problems that they emphasize, they have much in common: the reliance on experimental method; the anchoring of concepts to

² A parallel distinction between structure and function occurred in the history of biology with the separation of anatomy and physiology. Debates in the history of biology were not so much whether structural or anatomical problems could be distinguished from functional or physiological problems as whether one or the other of these problems should be given priority. The arguments rested upon how, in an organism, the structure of an organ might determine its functions and how, in the evolution of a species, the functions of an organ might determine its structure. Fortunately for biology, this distinction between structure and function did not become correlated with other distinctions, such as that between mechanism and vitalism (Catania, 1978).

experimental manipulations and observations; and the assumption that our subject matter, however complex, is orderly and not capricious.

IMPLICATIONS OF THE ANALYSIS OF BEHAVIOR

To assert that behavior is orderly is simply to say that we believe behavior can be understood. If human actions were capricious, we could not expect to discover any general principles by which to interpret these actions. Yet the assumption that our actions are caused or determined or predictable is sometimes regarded as incompatible with any dignified view of human nature. We value freedom, but freedom seems inconsistent with a scientific view of behavior. How can this difficulty be resolved?

We say that we are free when we are able to make choices. We choose to work or to play, to spend time with one companion or another, to accept things as they are or try to change them. How are any of these choices made? If the choices are important, they are not left to chance; we say that we weigh the alternatives. In each instance, the potential consequences of engaging in one alternative or the other influence what we choose. How can potential consequences enter into our decision, however, except by way of our earlier experience with other choices? It is in this sense that we may speak of our choices as determined. Would we wish to be free if by that we meant that our choices would be unaffected by their potential consequences, or that we were indifferent to the possible outcomes?

This point is sometimes misunderstood because we are often more concerned with another aspect of the concept of freedom. We are not so much worried about *whether* our choices are determined as we are about *how* our choices are determined. We say we are less free when our choices are limited, and we are then concerned with freedom and coercion. To the student, for example, there is an important difference between enrolling in a course because of one's interest in the subject matter and enrolling in a course because it is required; the choice to enroll may be as predictable or determined in one case as in the other, but we speak of the latter choice as forced, not free.

Once freedom has been defined in terms of the availability of choices, it even becomes possible to examine it experimentally. If an organism is said to be free only when alternatives are available, it can further be said to value freedom only if it prefers the availability of alternatives to their unavailability. Voss and Homzie (1970; see also Catania, 1975), for example, studied food-deprived rats' preferences for two paths leading to food. One path consisted of a single predetermined route; the other allowed the rat to choose between two subpaths. Fourteen of fifteen rats in this experiment preferred the path that allowed a choice of subpaths. Whether this preference had phylogenetic origins or was somehow acquired during each rat's lifetime, it can at least be concluded that the concept of freedom has a place in the analysis of behavior. And if a preference for freedom turns out to be part of the phylogenetic endowment even of the rat, then corresponding human

preferences take on a special significance, because they need no longer be attributed merely to the practices of particular human cultures.

We may prefer free choices, but we must also recognize that our choices are restricted by the nature of the physical world, by our own predispositions and capacities, and by the deliberate and the accidental actions of others. In our society, persuasion, legal sanction, threats, enticements, force, all are techniques that may be used to limit our choices. The techniques may be obvious or subtle, direct or devious, but they all enter into the potential consequences of our various alternatives and thereby influence our decisions. The issue becomes one not of whether our choices are determined, but rather of how our choices are manipulated, and by whom. One contribution of psychology may be simply to demonstrate the various ways in which our behavior can be influenced. If the determinants of human behavior can be identified, then these determinants must not be ignored; we cannot eliminate them simply by choosing not to study them. If our understanding of these determinants leads to the refinement of techniques to control human behavior, then this understanding must not be discouraged; we can best defend against the misuse of such techniques if we understand them.

That human behavior is determined or at least predictable is not, of course, an empirical fact that can be confirmed or disproved. It is an assumption with which a scientific account often proceeds. This does not imply that we expect to be able to interpret any instance of behavior. There are limits to what we can know. It is tempting to ask a psychologist to explain why someone behaved in a particular way, what led up to a particular incident, or how someone came to have particular interests, fears, or attachments. But the psychologist often has so little information available that only a plausible interpretation can be provided. This situation is not different except in degree from that in other sciences. Just as the principles of aerodynamics are not invalidated if we are unable to account for every twist and turn in the particular path of a particular falling leaf, the principles of behavior are not invalidated if we are unable to account for every detail of a particular organism's performance on a particular occasion. We must not expect too much.

OVERVIEW

We turn now to the science of behavior. If the study of learning is about how behavior can be modified, we must first consider what behavior is, how it can be investigated, and what vocabulary might best describe it. The next chapter considers some basic properties of behavior, some procedures that have been used in the study of learning, and some problems created by our ways of talking about behavior. Later chapters examine specific procedures in more detail. They begin with the simplest procedure, presenting a stimulus, and then move through the arrangement of consequences and the nature of behavioral units to more complex procedures, including those involving the signalling effects of stimuli. The principles developed are then applied to an analysis of language, which provides a transition to a treatment of human verbal learning and memory. The concluding topics briefly

examine behavioral processes called *cognitive*, including problem solving and thinking. We will find throughout that the most *profitable course is one that* stays close to the data; we will worry less about psychological theory than about *the* proper description of psychological findings. For example, it typically will be more useful to describe what an organism has learned or has remembered than to attempt to explain its learning or its memory. In our study of learning, it is important to recognize what remains out of our reach.

Behavior

BEHAVIORAL HIERARCHIES

STIMULI AND RESPONSES

OPERATIONS AND PROCESSES

The observation of behavior

The presentation of stimuli

Consequential operations

Stimulus-control operations

Stimulus-control operations and elicitation

Stimulus-control operations and consequences

RECAPITULATION

TWO

The word *behavior*, like *habit*, *inhibit* and *ability*, is related to the Latin *habere*, to hold, possess or have. The prefix *be-* became attached in such words as the Old and Middle English *behabban* and *behaven*, which meant to hold or bear oneself in some way. These words were closer to the sense of comportment or demeanor than to the more contemporary sense of activity, just as *habit* once referred more commonly to what was worn than to what was habitually done. A further aspect of this etymology or word history is that the Latin *habere* and the Latin *capere*, to take or seize, can be traced to closely related Indo-European roots. *Capere* leads to such psychologically significant words as *concept* and *perception*; thus, these words and *behavior* are distant relatives.

Behavior is no easier to define than learning. We may say glibly that behavior is anything an organism does. We soon wonder, however, whether this definition is too global. Should we count respiration, the beating of the heart or metabolism along with muscular movements and glandular secretions? We describe behavior with verbs: people walk, talk, think, do things. But although we may say that someone breathes, we are not likely to say that someone heart-beats. We also like to distinguish between active and passive actions. People bleed when cut and fall when they lack support, but we may not wish to speak of their bleeding or their falling as behavior.

Let us not try to resolve this problem. Our aim is to examine some properties of behavior, and to consider how these properties can be described. The phenomena of behavior are varied even though they sometimes share common names, so we shall probably do better by considering examples than by attempting definitions. We can deal with specific examples without much risk of misunderstanding.

When we observe an organism, all we see are properties of the organism's environment and properties of its behavior. We call these properties *stimuli* and *responses*, but neither a stimulus nor a response is of interest by itself. We study a stimulus to find out how it may affect responses, and we study a response to find out how it may be affected by stimuli. An experimental analysis determines what

kinds of relations exist between stimuli and responses, and how these relations can come about.

BEHAVIORAL HIERARCHIES

Observing the organism is the simplest thing we can do. For example, if we were interested in the behavior of a rat, we might just watch it (cf. Bindra, 1961). At first the rat might remain motionless. Eventually it might begin to move about. We might notice it sniffing, or licking and grooming itself, or urinating and defecating. Some changes in its behavior might ordinarily be inaccessible to us; only by setting up the necessary recording equipment could we note changes in its respiration, heart rate or blood pressure.

These and other responses could easily be inventoried. We could rank the various responses by how often they occurred, list the sequences in which they were ordered, and determine which could occur simultaneously and which only in succession. But to find out more, we would have to provide an environment in which the rat could engage in other different kinds of behavior.

Let us hypothetically construct such an environment by providing entrances into various compartments that give the rat access to different stimuli. Figure 2-1 shows some possibilities: compartments containing a full food hopper, a drinking tube filled with water, an activity wheel, or nesting material; a small dark compartment and a large bright compartment; compartments in which the rat's entry turns on a display of dim lights on the wall, quiet music, loud noises, a cold shower, or a blast of hot air; and, finally, a compartment with an electrified grid floor.

Here again we might observe locomotion, licking or grooming, urination or defecation. But we would also begin to learn about the rat's behavior with respect to these various stimuli. Presumably if the rat lived in this area for some time, it would begin to enter the food compartment occasionally and eat from the food hopper, or enter the water compartment and drink from the drinking tube, or enter the activity compartment and run in the activity wheel. After a few days, it might always sleep in the compartment with nesting material, or it might prefer the small dark compartment and perhaps would even carry some nesting material there. It probably would seldom stay for long in the large bright compartment. Eventually, it might spend brief periods looking or listening in the compartments with lights or music. But after only a few entries, it would probably rarely return to the compartments with loud noise, the cold shower, hot air, or the electrified grid.

We must assume that the rat's performance involves learning; the rat must learn where each compartment entrance leads. But our present concern is how the rat chooses to distribute its time among the compartments and, therefore, its tendency to engage in the various kinds of behavior that each compartment makes possible. Over any period of time, we could describe the rat's behavior in terms of a *hierarchy* of responses, an ordering based on the relative frequencies with which the various responses occur. The organization of behavior in terms of the relative likelihoods of different responses has been referred to as a habit-family hierarchy (Hull, 1943).




Figure 2-1. A hypothetical behavior space for the rat. When the rat is placed in the open area in the center, the order and relative frequency of visits to the various compartments can be recorded. On the basis of such a record, the rat's behavior can be described in terms of a response *hierarchy*.

During the night, when rats are most active, the rat might go to the food compartment and eat for a few minutes, then move on to the activity-wheel compartment and the water and briefly drink, and finally go to the activity-wheel compartment and run for a while. At first, we might say that eating is highest in the hierarchy, with drinking next and then running. Once some eating has occurred, drinking takes its place at the top of the hierarchy. After some water has been consumed, drinking is displaced by running. In other words, the order of responses in the hierarchy changes from time to time. In the daytime, the responses of eating and drinking and running all might be relatively low in the hierarchy; the rat might spend most of its time sleeping in the small dark compartment.

At this point, one might argue that instead of the hierarchical ordering of the rat's responses, we should discuss the rat's behavior in terms of its relative

preferences for different stimuli. We might then describe the rat's movement from eating to drinking to running simply by noting that food was at first the most preferred stimulus and that later this preference shifted to water and to the running wheel. In this language of preference, however, we would still be dealing with exactly the same events. Thus, this example only illustrates how closely connected are our concepts of behavior and environment: we define environmental events in terms of their effects on behavior, and we define behavior in terms of its interaction with the environment.

At any given time, of course, the rat can only be in one compartment. While the rat eats, for example, we can say little about the relative positions of listening to music or watching lights in this rat's response hierarchy. But we could find out about them simply by closing all the compartment entrances except those two and observing how the rat distributes its time between music and lights. What we learn about the rat still depends on what we choose to measure. Describing the rat's behavior in terms of response hierarchies will necessarily depend on what we make available to the rat in the compartments; some aspects of the rat's behavior will inevitably be omitted. For example, the compartments in our hypothetical apparatus will tell us nothing about the rat's social or sexual behavior. If we wished to study how this rat interacts with other rats, we would want to add more compartments, some containing a male rat, others containing a female rat, and still others containing different-sized groups of rats of one or both sexes.

Nevertheless, we could begin to classify the rat's responses according to the different probabilities with which the rat engages in them. The rat often enters the compartment where food sets the occasion for eating, whereas it rarely if ever enters the compartment where the electrified grid sets the occasion for jumping and squealing. The stimuli that generate these different responses have been called by different names; the former are sometimes referred to as *appetitive* or *reinforcing* stimuli, and the latter as *aversive* or *punishing* stimuli. The rat will produce the former by entering the appropriate compartment, but not the latter. Yet this classification is incomplete; the rat's failure to enter certain compartments is not evidence enough that the stimuli in those compartments are aversive. Some stimuli may be effectively *neutral*. For example, the rat may infrequently enter the compartments with music or with dim lights simply because the rat is usually indifferent to such stimuli; listening to or looking at such stimuli may be relatively low in the response hierarchy. Thus, we must be concerned not only with the rat's entry into different compartments, but also with its departures.

At this point, simply watching the rat is not enough; we must intervene. To assess the effects of the stimuli in those compartments that the rat rarely enters, we must place the rat in those compartments and measure how rapidly it leaves. We might find that the rat exits from the compartment with the electrified grid more rapidly than from those with cold showers or hot air, that it exits from these more rapidly than from the large bright compartment or the one with loud noise, and that it exits from these in turn more rapidly than from the compartments with music or dim lights. On the basis of these exits, we could order the stimuli of the various compartments. The ordering would be relative, however, because it is not obvious which compartment we should take as a baseline, or point of reference.

Even the open central area is a space where the rat may choose to spend some time. We must recognize that the rapidity of exit will depend not only on the stimuli in a compartment, but also on the availability of other compartments. For example, if all the other compartment entrances were closed off and the rat learned that it had to remain in the open central area, we might not expect it to leave the large bright compartment as quickly as it would if it learned that entry to all the other compartments were still available.

From such observations, we must conclude that stimuli do not fall neatly into a discrete classification. We cannot divide the environment into three simple classes of events, such as appetitive, neutral, and aversive. Instead, each stimulus, and the responses it occasions, must be evaluated relative to the other available stimuli. At any given time, the environment provides a continuum or range of possibilities: the organism may be highly likely to expose itself to some stimuli, may be relatively indifferent to others, and may not only fail to expose itself to but may also quickly remove itself from still others. The ordering of events along this continuum can change from one moment to another. Changes can result from earlier responding, as when eating varies as a function of the time since the last meal, or from changes in stimulus properties. If we are interested in the effects of music and vary its loudness, for example, we might find that the rat enters that compartment at quiet levels, is indifferent at moderate levels, and avoids the compartment at more intense levels. Thus, we would have to conclude that music is reinforcing, neutral or aversive, depending on its level.

We cannot say what the effects of a given stimulus will be except by observing behavior. We can call food appetitive or reinforcing only on the basis of behavior with respect to food; we can call dim lights neutral only on the basis of behavior with respect to dim lights; we can call electric shock aversive or punishing only on the basis of behavior with respect to electric shock. These examples all demonstrate that the characteristics of different stimuli cannot be specified independently of an organism's behavior; but neither can the characteristics of responses be specified independently of an organism's environment.

STIMULI AND RESPONSES

This hypothetical case dealt with both the stimuli and the rat's responses in each compartment. The close relation between these environmental and behavioral events was implicit. Let us now be more explicit about such relations, by turning to a different and more interesting organism, the human infant.

Again we begin simply by observing. We watch for a time, and we notice movements of the hands or arms or legs. Perhaps at some point the infant begins to cry. If the crying stops without our intervention, the infant might sleep or might lie quietly with its eyes open. If we look closely, we might see the eyes moving, although it would be difficult to judge just what the infant was looking at, or whether it was looking at all. We could begin to catalogue the various things that the infant did, and we might discover that particular movements usually occur in particular sequences. Just as with the rat, we would not learn very much if we only

watched; we could not say much more than that particular movements occurred more or less often and more or less in particular orders.

Again, we need not be restricted simply to watching. We can do things to the infant. We might touch it, rock it, move objects before its eyes, make sounds, or place a nipple in its mouth. We should not be surprised if the infant responds to each of these events in a characteristic way. If the touch were to the palm of the hand, for example, the infant would probably clench its fist, grasping the object that touched it. The vocabulary for these events is already familiar: we call the touch to the palm a *stimulus*, and the grasping a *response*. A particular environmental event is followed by a particular instance of behavior. In this case, we are not interested in the stimulus alone or in the response alone; we are interested in their relation to each other. We call this relation, the reliable production of a particular response by a particular stimulus, a *reflex*.

We will consider the reflex in chapter 3. The important point here is that the term *reflex* is simply a name for a behavioral relation: an observed correlation between a particular stimulus and a particular response (Skinner, 1931). The concept is neither a theory nor an explanation. We observe a particular relation between environment and organism, and we name it. If our naming has been effective, we will simply know whether or not, with respect to other relations we might observe, the same name is appropriate.

Let us return to the infant. Besides the production of grasping by a touch to the palm, we might wish to catalog other examples of reflexes: the production of crying by a loud noise, or of sucking by a nipple in the mouth, or of blinking or pupillary contraction by a flash of light. These are not the only kinds of relations possible, however. When stimuli produce responses, the environment is acting on the infant; but the infant can also act on the environment. Its crying, for example, often brings its mother's attention. Crying, then, is a response that can produce an environmental consequence: the presence of the mother. This behavioral relation involves stimuli and responses, but we cannot call it reflex because the stimuli and the responses do not occur in the appropriate order.

The relations can become more complicated. If the infant moves its eyes while the lights are on, the eye movement may change what the infant sees. Eye movements cannot have this effect when the lights are off. Thus, the infant may come to look about in the light but not in the dark. In the presence of one stimulus, the light, moving the eyes may have consequences; it may produce another stimulus, some new thing seen. These consequences cannot occur in the dark. Thus, the relation involves three terms: a prior stimulus, the light; a response, eye movement, in the presence of this stimulus; and a subsequent stimulus, what is newly seen as a consequence of this response. This three-term relation, prior-stimulus—response—subsequent-stimulus, is important because an organism's behavior depends both on the situations in which it finds itself and on the consequences of its responses in these situations. To deal with the kind of relations that can exist between stimuli and responses, however, we must recognize another problem: what kinds of events are these?

A stimulus is an environmental event, but such events have varying degrees of complexity. In our example, we regarded the infant's mother as a stimulus. The

infant's environment is certainly different when its mother is present than when she is absent. Yet what sort of stimulus is she? Early in its life, the infant cannot recognize her face or distinguish her from others. We do not know which aspects of her looks, her voice or her touch are important to the infant. We might wonder whether the infant would react to her any differently than usual if she approached wearing a surgical mask. Nevertheless, despite our ignorance with respect to these questions, we do not doubt that the mother is an important feature of the infant's environment. Lacking a more detailed account of the specific features of the mother to which the infant may respond, we may still find it useful to speak of the effects that she has as she comes and goes in the infant's world.

This example illustrates two different problems. If we try to analyze the visual, auditory and tactile features of the mother that are important to the infant, we become concerned with the structure of this complex stimulus, the mother. We might ask how it is that the infant will be able to learn to respond to a particular individual as mother despite changes in her dress or hair style, her facial expression or posture. But if we concentrate instead on how the mother interacts with the infant's responses, we are concerned with the functional significance of the mother in the infant's environment. For example, if an infant was crying, we would not have to know whether the infant recognized its mother by her face, her hair or her voice to judge whether the mother's presence made a difference in the crying; it would be enough to note that when the mother went to the infant the crying stopped. Later, we will typically be interested in simpler stimuli: lights, sounds, food in the mouth. But even with simpler stimuli we will have to remember that structural problems, in analyzing stimulus properties, are different from functional problems, in analyzing the interaction between stimuli and responses.

And what about responses? How shall they be dealt with? In describing responses, we encounter at least two difficulties. The first is that behavior is not repeated exactly from one instance to the next. If the infant grasps an object on two different occasions, the grasping will not be the same each time. The difference *may be small, in the force of the grasp*, for example, or in the exact placement of the fingers. But if *there is any difference at all, we must worry* whether we should regard the two grasps as two instances of the *same response or as two different responses*. We must speak not of individual responses, but of classes of responses having common properties.

The second difficulty is that responses are sometimes adequately described in terms of movements, but at other times the description must include the environment in which the responses occur. For example, suppose we wish to compare the infant's grasping of an object with its clenching of its fist. In terms of the muscles that move, grasping an object with the right hand and clenching the right fist have more in common than grasping an object with the right hand and grasping an object with the left hand. Yet it may sometimes be more important to speak of the act of grasping an object, no matter which hand the infant uses, than to speak of the movement of closing a particular hand. An account of behavior must distinguish between movements, responses defined by their form or the musculature used, and actions, responses defined by their relations to the environment. Consider how often we speak of doing things, going places, or manipulating

objects, without regard to the particular manner of performing these actions.

Even in the absence of movement, we may sometimes conclude that behavior has occurred. For example, an infant will typically grasp an adult's fingers sufficiently tightly that the infant can be lifted into the air. Once lifted, the infant may not move while holding on, and yet the very fact that the infant does not fall leads us to conclude that the grasping response continues. Similarly, if we see an adult standing still, our judgment that the adult is behaving stems partly from our knowing that the adult would fall down if unconscious or dead. It might be argued that the standing adult is in fact moving in small, unnoticed ways, but even if minor postural adjustments occur in standing still, we do not have to observe them to conclude that the adult is behaving. Thus, not all movements need to be instances of behavior, and not all instances of behavior need be movements. Whether behavior involves movement or not, it is important to examine its consequences. Thus, the critical feature of the infant's grasping and the adult's standing is simply that these responses have an effect; they determine whether or not the organism falls.

OPERATIONS AND PROCESSES

One more distinction remains before we consider some phenomena of learning: the distinction between *operations* and *processes*. We study the relation between events in the environment and the organism's responses by manipulating the environment and observing how this affects what the organism does. We operate on the organism's environment or, in other words, we perform experimental operations. In the analysis of behavior, operations are *what the experimenter does or arranges*, and processes are *the changes in behavior that result*.¹ Particular learning procedures can be described in terms of these operations, taken either alone or in some combination.

The simplest operation, of course, is merely to *observe behavior*. The behavior that we observe tells us what an organism is capable of doing in such circumstances. But because we have no control over events when we only observe, we may be unable to draw conclusions about the causes of behavior. We therefore must intervene, and the simplest intervention is to *present stimuli*. When a stimulus produces responses, we may speak of the process of *elicitation*. A more complicated intervention is that of arranging the environment so that the organism's behavior has *consequences*. Once responses have consequences, they may occur more or less often, and therefore consequential operations lead to the processes sometimes called *reinforcement* and *punishment*. Finally, we can arrange that stimuli *signal* events, such as the presentation of other stimuli or the availability of particular consequences. We then speak of *stimulus-control* operations; these operations can only occur in combination with one of the simpler operations, presenting stimuli or arranging consequences. Behavior may then depend on whether the signalling stimulus is present or absent. Thus, the basic operations that

¹ A medical analogy makes this easier to remember: a surgical operation is what the doctor does to the patient, and processes are the physiological effects of the operation, such as changes in circulation, metabolism, respiration, and so on.

we will consider are: *observing an organism's behavior, presenting stimuli, arranging consequences for responses, and signalling either stimuli or consequences.*

The observation of behavior

What must we do to observe behavior? When we considered the rat, we recognized that its behavior would be interesting only if we gave it an interesting environment. Instead we might logically move in the other direction, avoiding any contamination of the rat's behavior by environmental interventions. We could fit the rat with goggles to exclude visual stimuli and with ear plugs to exclude sounds and we could remove odors with a ventilating system. We might realize that the rat could still touch things, including parts of its own body, and so we might then arrange a suit of hollow tubes that held the rat's legs so that tactile contact was reduced for at least its paws. This might still not satisfy us, because the force of gravity would produce pressures at those places where the suit met parts of the rat's body and would allow spatial orientation. The next step would be to send the suited rat up in a space shuttle, where gravity could be eliminated. Yet even if we accomplished this much, what could we say about the rat's behavior? What might we observe the rat doing?

The rat example is hypothetical, but experiments on sensory deprivation have placed humans in environments that approximate the conditions of minimal stimulation suggested for the rat (e.g., P. Solomon, Kubzansky, Leiderman, Mendelson, Trumbull, and Wexler, 1961). The problem is that in environments like these, for human as well as for rat, there is not much to do; there is no place to go and no one to see or to speak to. Although humans in such environments report a range of activities during their waking time, from thinking to hallucinating, it is not surprising that they spend a fair proportion of time sleeping.

So, we might conclude that we had been right in the first place. To observe interesting behavior, we must place the organism in an interesting environment. As an example, consider Köhler's experiments on chimpanzee behavior (Köhler, 1927). From 1913 to 1917, Köhler observed chimpanzees maintained at the Anthropoid Station on Tenerife, an island northwest of Africa. In some of these experiments, bananas or oranges were placed in visible but inaccessible locations, but materials within the chimpanzees' chamber could be used as tools to gather in the fruit. The following describes the behavior of the male chimpanzee, Sultan:

The six young animals of the station colony were enclosed in a room with perfectly smooth walls, whose roof—about two metres in height—they could not reach. A wooden box . . . , open on one side, was standing about in the middle of the room, the one open side vertical, and in plain sight. The objective was nailed to the roof in a corner, about two and a half metres distant from the box. All six apes vainly endeavored to reach the fruit by leaping up from the ground. Sultan soon relinquished this attempt, paced restlessly up and down, suddenly stood still in front of the box, seized it, tipped it hastily straight towards the objective, but began to climb upon it at a (horizontal) distance of half a metre, and springing upwards with all his force, tore down the banana. About five minutes had elapsed since the fastening of the fruit; from the momentary pause before the box to the first bite

into the banana, only a few seconds had elapsed, a perfectly continuous action after the first hesitation (Köhler, 1927, pp. 39-40).

In many instances, of course, chimpanzees made unsuccessful attempts to solve problems such as these; such attempts were aptly described as *fruitless*.

Köhler was concerned with the ways in which chimpanzees could respond to the relations among objects and events in the environment, and he discussed these and related observations in terms of the chimpanzee's intelligence and insight. His major contribution was in demonstrating what chimpanzees were capable of doing, and the readers of his time were impressed by his descriptions of chimpanzee performance. The problem was that it was not possible to say from observation alone where the behavior came from. Was Sultan able to solve a particular problem because of some inherited cognitive disposition, or because of features that the problem had in common with some situation that he had already encountered, or because of exposure to the behavior of other chimpanzees that he was able to imitate, or because of some combination of these and other factors? The term *insight* seemed appropriate because of the suddenness with which a solution to a problem often emerged, and problem solving that seemed insightful led to further questions: whether learning took place abruptly or gradually, and whether this type of problem solving could be explicitly taught. Debates about the instructability of nonhuman primates continue to the present. But such issues cannot be resolved from observation alone. Observation by itself only rarely can indicate the sources of behavior.

Strictly, however, Köhler did more than simply observe behavior. He arranged environments within which to make his observations. Observation without intervention is difficult to achieve, because even bringing an organism into captivity is itself an intervention. Although he did not do so, Köhler instead could have searched out wild chimpanzees to observe in their natural habitats, or could have released some captive chimpanzees and attempted to track them in the wild. But the problem of keeping up with them as they moved on the ground or through the trees would have been compounded by the problem of also keeping hidden from them. To observe successfully in the wild, one must know the possible effects of human sight and scent and sound on the organism under observation (e.g., van Lawick-Goodall, 1968). To study these effects, one must present appropriate stimuli. Presenting stimuli, therefore, is a virtually inevitable component of any study of behavior.

The presentation of stimuli

We have seen that Köhler did in fact present stimuli, by arranging environments for the chimpanzees he observed. Let us turn now to a different set of examples. The following, by the ethologist Tinbergen, describes the first feeding of the Herring Gull chick:

Sometimes the parent stands up and looks down into the nest, and then we may see the first begging behavior of the young. They do not lose time in contemplating or studying the parent, whose head they see for the first time,

but begin to peck at its bill-tip right away, with repeated, quick, and relatively well-aimed darts of their tiny bills. They usually spread their wings and utter a faint squeaking sound. The old bird cannot resist this, and if only the chicks persist it will feed them. First the parent stretches its neck, and soon a swelling appears at its base. It travels upward, causing the most appalling deformations and the most peculiar turnings and twistings of the neck. All at once the parent bends its head down and regurgitates an enormous lump of half-digested food. This is dropped, and a small piece is now picked up again and presented to the chicks. These redouble their efforts, and soon get hold of the food, whereupon the parent presents them with a new morsel. Now and then the chicks peck at the food on the ground, but more often they aim at the parent's bill, and although this aiming is not always correct, it rarely takes them more than three or four attempts until they score a hit. (Tinbergen, 1960, p. 178)

So far, we have here only some observations of chick behavior. But they involve the effects of stimuli, and therefore prompt some questions. What exactly are the critical features of these special stimuli presented by the parent gull? Are some more important than others? Are they the most effective ones possible?

Tinbergen set out to answer these questions by preparing stimuli that resembled the parent gull in various ways. He then compared the pecking generated when these stimuli were presented to recently hatched Herring Gull chicks. The Herring Gull parent has a white head and a yellow beak with a red patch near its tip. A beak with a black or a blue or a white patch produced less pecking than one with a red patch, but a beak with a patch of any color produced more pecking than a beak without any patch at all. Compared to the red patch, the color of the beak and of the head were relatively unimportant in generating pecking. In fact, if the model did have a beak with a red patch on it, even the presence or absence of a head made little difference.

Tinbergen also varied the shape of the beak, as illustrated in Figure 2-2. Next to each stimulus, pecking is shown as a percentage of the number of pecks generated by the normal beak shape, at the top (100%). The red patch and other color differences were eliminated, so that changes in pecking could be attributed solely to changes in shape. The interesting feature of varying shape is that, although most models produced less pecking than the model with the normal beak shape, the model with an elongated beak (bottom) produced considerably more pecking than any of the others, including the one with the normal beak shape. Because of its effectiveness relative to the normal shape, Tinbergen referred to this model as a *supernormal* stimulus, but he also speculated that the shape of this particular model might be more like what the chick first sees of its parent's beak, as it looks up from beneath the parent's head, than the profile view used for most of the other models.

Tinbergen also varied other properties of the Herring Gull beak, such as movement, slant, and height above the ground. His analysis enabled him to construct a truly supernormal stimulus, a red pencil-shaped rod with three narrow white bands that generated more pecking than an accurate three-dimensional model of a Herring Gull head. More important, he was able to specify those features of the parent's head that were important in generating pecks, and to distinguish them from unimportant features. In other words, by presenting stimuli and observing their effects,




Figure 2-2. A series of models used in analyzing which properties of the parent gull's beak produce begging pecks in the newly hatched Herring Gull chick. Pecks to each other model are expressed as a percentage of the reference level (100%) provided by pecks to the top model. (Adapted from Tinbergen and Perdeck, 1950, Figure 15.)

Tinbergen was able to identify the critical *structure* of the stimuli that generated pecking in the newly hatched Herring Gull chick.

Stimulus presentations are a common feature of research conducted by ethologists, whose concern is the evolution of species-specific behavior patterns and behavior in an organism's natural habitats. One effect of stimulus presentations, as we have just seen, is the production of responses. This process, an outcome of presenting stimuli, is called *elicitation*; the stimulus is said to elicit a response. In the language of ethology, the critical stimuli or stimulus features are called *releasers*, and the behavior they produce is called a *fixed action pattern*. But variations in vocabulary should not obscure the simplicity of the basic operation of presenting stimuli.

By presenting stimuli, we can determine what features are critical in generating a particular instance of behavior, but we cannot necessarily deal with how this behavior changes over time. Data from the Laughing Gull chick provide an example. Feeding in the Laughing Gull chick differs in detail from that of the Herring Gull, but also includes a begging peck at the parent's beak followed by the parent's regurgitation of partly digested food that the chick then eats. The accuracy of the begging peck was tested by presenting beak models to chicks of various ages. Only about one-third of the pecks of newly hatched chicks struck the model, as opposed

to more than three-quarters of the pecks of two-day-old chicks (Hailman, 1969). Was the improved accuracy a function of changes in coordination or in visual experience or in other factors?

One possibility is that the consequences of accurately aimed pecks differ from those of poorly aimed pecks. In the Laughing Gull's natural habitat, an accurately aimed peck is more likely than a poorly aimed peck to be followed by the parent's regurgitation of food. Even without this outcome, an accurately aimed peck will have a different consequence from a poorly aimed peck, because only the accurately aimed peck will contact the parent's beak. If contacting the parent's beak and then feeding are important consequences for the Laughing Gull chick, we might assume that responses producing these consequences will become more frequent than those that do not or, in other words, that accurately aimed pecks will increase relative to poorly aimed pecks. Hailman's observations in fact suggest that the chick's behavior is determined by such consequences:

If an inexperienced chick is too close to the target at first, its pecking thrust against the bill or model is so strong that the chick is thrown backward as much as an inch. If the chick starts out too far from the target, the pecking thrust misses and the chick falls forward as much as two inches. Older chicks rarely make such gross errors, suggesting that the experience of overshoots and undershoots has helped the chick learn to adjust its distance. (Hailman, 1969, p. 100)

To study such effects, it is not enough simply to present stimuli. A more complex operation must be arranged: Stimuli must be presented as consequences of the organism's behavior.

Consequential operations

Again we move to a new set of examples, this time from Thorndike's research on animal intelligence (Thorndike, 1898). The critical difference between Thorndike's research and Köhler's was that Thorndike systematically observed changes in behavior over many repetitions of an organism's behavior in a given setting, rather than restricting his attention to single instances of a problem solution. The problems he studied did not lend themselves to the sudden or insightful solutions seen by Köhler, but they did show that consequences can be important determinants of behavior. Thorndike showed that responding in a given situation depends significantly on the consequences such responding has had in the past.

Thorndike described his experimental procedures in this way:

I chose for my general method one which, simple as it is, possesses several other marked advantages besides those which accompany experiment of any sort. It was merely to put animals when hungry in enclosures from which they could escape by some simple act, such as pulling at a loop of cord, pressing a lever, or stepping on a platform. . . . The animal was put in the enclosure, food was left outside in sight, and his actions observed. Besides recording his general behavior, special notice was taken of how he succeeded in doing the necessary act (in case he did succeed), and a record was kept of

the time that he was in the box before performing the successful pull, or clawing, or bite. . . . If, on the other hand, after a certain time the animal did not succeed, he was taken out, but *not fed*. Thorndike, 1898, pp. 5-6)

One of Thorndike's puzzle-boxes is illustrated in Figure 2-3. In such apparatus, Thorndike studied cats, dogs, and chicks. The following description was offered as characteristic of the behavior of most cats:

When put into the box the cat would show evident signs of discomfort and of an impulse to escape from confinement. It tries to squeeze through any opening; it claws and bites at the bars or wire; it thrusts its paws out through any opening and claws at everything it reaches; it continues its efforts when it strikes anything loose and shaky; it may claw at things within the box. . . . The cat that is clawing all over the box in her impulsive struggle will probably claw the string or loop or button so as to open the door. And gradually all the other non-successful impulses will be stamped out and the particular impulse leading to the successful act will be stamped in by the resulting pleasure, until, after many trials, the cat will, when put in the box, immediately claw the button or loop in a definite way. (Thorndike, 1898, p. 13)

As a consequence of its responding, the cat escaped from confinement and also gained access to food. We may assume that both escape and food were important in making the successful response gradually dominate over other, unsuccessful ones. In either case, the procedure cannot be reduced simply to the presentation of stimuli. Once the cat got out of the box, it confronted new parts of the environment, including the food it found there. These features of the environment were not merely presented to the cat; they became available to the cat as a




Figure 2-3. A Thorndike puzzle-box (Thorndike, 1898, Figure 1). In most situations that Thorndike studied, the animal had only a single way to open the door. In the box shown, three different methods for opening the door are illustrated: a treadle inside the box (A); a wire or string that can be reached from inside the box (B); and two outside latches that can be reached from within to free the door (C). The door (D) was usually counter-weighted so that it opened by itself once the animal made the required response.

consequence of what it did. The cat's behavior changed as a result of this consequential operation, which made events depend on the cat's behavior. Contemporary experiments, which often examine simpler responses in simpler situations, are similarly concerned with the relation between responses and their consequences.

We arrange consequences for responses by constructing environments in which these responses can have specified effects. If we place food in the goalbox of a maze or a runway, for example, we create an environment in which a rat's movement from the startbox to the goalbox will have food as a consequence. After the rat has reached the food on one occasion, we may examine how the food as a consequence affects the rat's behavior by observing the rat's locomotion the next time we place it in the startbox.

In consequential operations, varied consequences can be arranged, from stimuli of obvious biological significance such as food and water to relatively minor changes in things seen or heard or touched. We need not even restrict ourselves to consequences that involve the production of stimuli. Responses can remove stimuli, as when operating a switch turns off a light. Responses can prevent stimuli, as when unplugging a lamp before repairing it eliminates the possibility of a shock. Responses can even change the consequences of other responses, as when replacing a burned-out light bulb makes the previously ineffective response of operating the light switch effective again. Each of these types of consequences can affect subsequent behavior. Nevertheless, presenting stimuli and arranging consequences are not the only experimental operations in the analysis of behavior.

Stimulus-control operations

The presentation of stimuli and the arrangement of consequences are seldom independent of other events in the organism's world. A flash of lightning typically precedes a clap of thunder. A traffic light typically signals the possible consequences of proceeding or stopping when one drives up to an intersection. These two examples illustrate how the signalling or discriminative effects of stimuli may be combined with stimulus presentations or with consequential operations. The behavior of organisms does not occur indiscriminately; behavior that occurs in one circumstance need not occur in another.

Stimulus-control operations and elicitation. The effects of stimuli that signalled the presentation of other stimuli were the basis for Pavlov's studies of conditional or conditioned reflexes (Pavlov, 1927). Pavlov's experiments were on the ways in which stimuli could acquire signalling properties. He examined responses to stimuli such as food, and then explored how these responses could be controlled by other stimuli presented in some temporal relation to food. Pavlov spoke of the effects of food in a dog's mouth in terms of the alimentary reflex, the components of which included both the glandular response of salivating and motor responses such as chewing and swallowing. He chose to concentrate on salivation in his studies because the available technology made salivating easier to measure than motor responses: by surgery, the duct of one of the dog's salivary glands was brought to the outside of the dog's cheek, where it was connected to a fluid system that allowed drops of saliva produced by this gland to be counted.

For one dog, the presentation of food was consistently preceded by the sound of a metronome. Pavlov gave the following account of the conditions necessary to establish a stimulus as a signal:

On several occasions this animal had been stimulated by the sound of the metronome and immediately presented with food—*i.e.*, a stimulus which was neutral of itself had been superimposed upon the action of the inborn alimentary reflex. We observed that, after several repetitions of the combined stimulation, the sounds from the metronome had acquired the property of stimulating salivary secretion and of evoking the motor reactions characteristic of the alimentary reflex. . . . Hence a first and most essential requisite for the formation of a new conditioned reflex lies in a coincidence in time of the action of any previously neutral stimulus with some definite unconditioned stimulus. Further, it is not enough that there should be overlapping between the two stimuli; it is also and equally necessary that the conditioned stimulus should begin to operate before the unconditioned stimulus comes into action. If this order is reversed, the unconditioned stimulus being applied first and the neutral stimulus second, the conditioned reflex cannot be established at all. (Pavlov, 1927, pp. 26-27)

Pavlov's conditioning experiments are effective demonstrations of how a stimulus-control or signalling operation can be superimposed on the simpler operation of stimulus presentation.

Stimulus-control operations and consequences. Instead of signalling the presentation of stimuli, a stimulus may signal occasions on which responses will have particular consequences: a stimulus-control operation can be superimposed on a consequential operation. For example, an environment can be arranged in which a rat's lever presses produce food when a light is on but not when the light is off (cf. Skinner, 1933, 1934a). In this instance, the light signals the current consequence of pressing the lever, and we might expect the rat to come to press the lever more often when the light is on than when it is off; we may say that the lever press is reinforced in the presence but not the absence of the light, and that the light is a discriminative stimulus.

The relations between a discriminative stimulus and the consequences of responding are elaborated in the following passage by Skinner (the term *operant* refers to a class of responses having particular consequences, and the term *reinforcement* refers to these consequences):

. . . the operant must *operate* upon nature to produce its reinforcement. Although the response is free to come out in a very large number of stimulating situations, it will be effective in producing a reinforcement only in a small part of them. The favorable situation is usually marked in some way, and the organism. . . . comes to respond whenever a stimulus is present which has been present upon the occasion of a previous reinforcement and not to respond otherwise. The prior stimulus does not elicit the response; it merely sets the *occasion* upon which the response will be reinforced. . . . Three terms must therefore be considered: a prior discriminative stimulus (S^D), the response (R^O), and the reinforcing stimulus (S^I). Their relation may be stated as follows: only in the presence of S^D is R^O followed by S^I . A convenient example is the elementary behavior of making contact with specific parts of

the stimulating environment. A certain movement of my arm (RO) is reinforced by tactual stimulation from a pencil lying on my desk (S^1). The movement is not always reinforced because the pencil is not always there. By virtue of the visual stimulation from the pencil (S^D) I make the required movement only when it will be reinforced. The part played by the visual stimulus is shown by considering the same case in a dark room. At one time I reach and touch a pencil, at another time I reach and do not. . . . In neither the light nor the dark does the pencil *elicit* my response (as a shock elicits flexion), but in the light it sets the occasion upon which a response will be reinforced. (Skinner, 1938, p. 178)

Skinner explored this three-term relation experimentally with lights as discriminative stimuli, the lever presses of rats as responses, and food pellets as reinforcing consequences. His example of reaching for and touching objects in the seen environment illustrates the broad range of situations to which the concept of stimulus control applies. The three-term relation, discriminative-stimulus—response—consequence, will be a recurrent theme. Each term is critical, because their combination distinguishes this case from other, simpler behavioral relations. In the Pavlovian situation, for example, in which a stimulus presentation is signalled, the organism's behavior has no effect on the sequence of events; no consequences are arranged for responses. Consider the earlier examples. Our blinking or startling at the lightning flash will not prevent the subsequent clap of thunder. But if a traffic light is red as we approach an intersection, our stepping on the brakes is occasioned by this stimulus only because we have learned the potential consequences of doing or not doing so. Only the second of these two examples involves all of the terms of Skinner's three-term contingency. An important difference in vocabulary accompanies these distinctions. *When a stimulus is the primary cause of a response, we say that the stimulus produces the response or that the response is elicited. But when a response occurs in the presence of a stimulus because the stimulus signals some consequence of responding, we say that the stimulus occasions or evokes the response and that the response is emitted.*

The development of stimulus control over consequential responding has been called discrimination learning. Early animal experiments often were concerned not so much with the nature of discrimination learning as with the sensory capacities of organisms (Yerkes, 1907). For example, rodent vision was studied by arranging two paths only one of which led to food (Yerkes and Watson, 1911). At the point where a rat had to choose the left or the right path, two stimuli were presented (e.g., a black card and a white card). The path to food varied from left to right, but was always indicated by the same card (e.g., black). Once the rat learned to take the path indicated by the stimulus correlated with food, the limits of its vision could be studied by substituting more similar stimuli (e.g., light and dark gray cards) for the original pair. Such experiments were laborious. To demonstrate discrimination learning might take hundreds of trials, if the rat learned at all. Several problems existed in this type of study, not the least of which was ensuring that the rat looked at the stimuli when it reached the choice point.

Apparatus improved over time. Figure 2-4, for example, shows the jumping-stand developed by Lashley (1930). Lashley described its advantages:

... it requires the animal to jump against the stimulus patterns from a distance, instead of to run past them. . . . I have usually trained the animals by placing the stand against the screen and allowing the animals to step through the open holes to the platform, then gradually withdrawing the stand until, in ten or fifteen trials, the distance of 25 cm. is reached. Cards are then placed in position and training in discrimination begun. (Lashley, 1930, pp. 454-457).

In Lashley's apparatus, rats typically learned to discriminate black from white with perfect accuracy within four or five trials, and even more difficult discriminations such as vertical versus horizontal could usually be mastered within less than fifty trials.

These are still cases in which discriminative stimuli signal consequences of responding, but they are more complex than the example in which a rat's lever presses produced food in the presence but not in the absence of light. There, our concern was only with how often lever presses occurred when the light was on and when it was off. Consider, however, the jumping-stand. We may regard it as involving only two responses and their respective consequences: jumping toward vertical and finding food, and jumping toward horizontal and landing in the net. But assuming that the left and right sides are irrelevant to the rat is only our prejudice. Perhaps we should treat it as involving four responses, each with a




Figure 2-4. The Lashley jumping-stand (Lashley, 1930, Figure 1). A rat was trained to jump from the stand (S) to one of the two doors (L and R). If the rat jumped to the correct door, the door gave way and the rat reached the food platform (FP). If the rat jumped to the incorrect door, the door remained fixed and the rat fell into the net below (N). The projecting metal sheet (M) prevented the rat from jumping too high. In the illustration, the right door (R) would be correct for a rat being trained to jump vertical lines.

particular consequence: jumping to vertical on the left, to vertical on the right, to horizontal on the left, and to horizontal on the right. Rats may be more likely to respond consistently on the basis of position, left or right, than on the basis of stimulus cards. For example, if the first four trials of vertical-horizontal training were set up with vertical on the right, as in Figure 2-4, we should not be surprised if on the fifth trial, with vertical on the left for the first time, the rat jumped right, toward horizontal. Until the fifth trial, jumping to the right was as effective in leading to food as jumping toward vertical.

The kind of discrimination in which a single stimulus is present or absent, as in the lever-pressing example, is called *successive* discrimination. That in which two or more stimuli are present at the same time and in which each is correlated with a different response, as in the jumping-stand example, is called *simultaneous* discrimination. Both illustrate stimulus-control operations superimposed on consequences of responding. The comparison between successive and simultaneous discriminations shows that stimulus-control operations come in varying degrees of complexity. In fact, an important feature of many experiments on behavior is the analysis of complex situations into their component operations.

RECAPITULATION

The study of behavior is concerned with relations between environmental events, *stimuli*, and the organism's actions, *responses*. These relations can be examined by analyzing how manipulations of the environment, experimental *operations*, produce changes in responding, behavioral *processes*. A critical first step is *observation* of behavior. Because we always observe behavior in some environment, it is typically not enough simply to watch an organism. To understand behavior, we must intervene by changing the environment. We can describe environmental changes in terms of three classes of experimental operations: *stimulus-presentation* operations, *consequential* operations, and *stimulus-control* operations. Table 2-1 summarizes these operations.

Presenting stimuli to an organism is the simplest operation. When we present stimuli, we can observe the responses produced. The major effect of the stimulus-presentation operation is to make particular responses more or less likely. For example, we might make a rat salivate by presenting food, but the rat would stop salivating and instead jump and squeal if we presented electric shock. Stimulus presentations are relevant to learning because the responses produced by stimuli can vary depending upon the conditions under which stimuli are presented.

The things that happen in an organism's environment do not always happen by themselves. Sometimes the organism makes them happen; behavior can have consequences. Environments can be arranged in which the organism's responses have effects. Arranging the environment so that an organism's responses produce or remove or prevent stimuli is called a consequential operation. For example, we could present food to a rat whenever it reared up on its hind legs, or electric shock whenever it moved to a certain floor area. As a result of consequential operations, some responses may occur more often and others less often. If an organism

TABLE 2-1

Basic Behavioral Operations.

<i>Operation</i>	<i>Description</i>	<i>Examples</i>	<i>Usage</i>
Stimulus-presentation operation	Stimulus <i>A</i> is presented	<i>Loud noise (A)</i> startles child. <i>Doctor shines light (A)</i> into patient's eye during exam.	Stimulus may <i>elicit</i> a response; response may <i>be elicited by</i> stimulus.
Consequential operation	Response <i>B</i> has Consequence <i>C</i> (e.g., a stimulus is produced or is terminated).	<i>Putting coin in vending machine (B)</i> produces <i>soft drink (C)</i> . <i>Touching hot stove (B)</i> produces <i>burn (C)</i> . <i>Light goes out (C)</i> when <i>switch is thrown (B)</i> .	Response may <i>be emitted</i> .
Stimulus-control operation			
Superimposed on stimulus presentation	Stimulus <i>D</i> signals presentation of Stimulus <i>E</i> .	<i>Lightning (D)</i> precedes <i>thunder (E)</i> .	Stimulus may <i>elicit</i> a response; response may <i>be elicited by</i> stimulus.
Superimposed on consequences	Stimulus <i>F</i> signals that Response <i>G</i> will have Consequence <i>H</i> .	<i>Red traffic light (F)</i> signals that <i>driving through intersection (G)</i> may lead to <i>accident or traffic ticket (H)</i> . <i>Ringling telephone (F)</i> signals that <i>answering (G)</i> may provide <i>opportunity for conversation (H)</i> .	Stimulus may <i>occasion</i> or <i>evoke</i> response; response may <i>be emitted in</i> presence of stimulus.

responds more often because its response changes its environment, we say that the response was rewarded or *reinforced*; if an organism responds less often for the same reason, we say that the response was suppressed or *punished*. The effects of consequential operations, the ways in which they change behavior, are critically important; we must deal with these operations when we consider how organisms can alter their world.

We can complicate matters by superimposing still another operation on stimulus-presentation operations or on consequential operations. We can arrange that these simpler operations are in effect only when some signal is present. Organisms do not behave indiscriminately. They do some things in some circumstances and other things in other circumstances. These relations are studied by arranging stimulus-control operations: one stimulus is used to warn of the imminent occurrence of another, or to signal the conditions under which a response may have some consequence. For example, we might present food to a rat only when a buzzer sounds, or electric shock only when a light is on. In these instances, a stimulus-control operation is superimposed upon a stimulus-presentation operation: the buzzer signals food presentations and the light warns of shock presentations. Stimulus-control operations can also be arranged in conjunction with consequential operations. For example, a rat's rearing up on its hind legs might produce food only if a lamp is lit, or its movement to a certain area of the floor might produce electric shock only if a tone is sounding. In these instances, a stimulus-control operation is superimposed upon a consequential operation: the lamp indicates when one response will produce food and the tone indicates when another response will produce shock. These stimulus-control operations may lead the changes in behavior produced by stimulus-presentation operations or consequential operations to occur only in the presence of the corresponding stimuli. When a stimulus signals an organism that some event is about to occur or that its behavior may have certain consequences, the organism may come to respond differently when the stimulus is present than when it is absent. This behavioral process is called *discrimination*, and the organism's behavior is said to be *under the control* of the discriminative stimulus.

Behavior is complicated because different stimuli can have different effects on different responses, and different responses can have different consequences. Nevertheless, these relations can be classified, studied, and distinguished from one another. Various learning procedures can be treated as combinations of these three basic types of experimental operations: *stimulus presentation* operations; *consequential* operations, such as reinforcement and punishment, in which these responses act on the environment; and *stimulus-control* operations, in which these other operations are signalled by discriminative stimuli. Alone or in combination, these operations provide a consistent means for analyzing learning procedures.

Elicitation

THE REFLEX

PROPERTIES OF ELICITED BEHAVIOR

Eliciting stimuli and response probabilities

Probabilities or relative frequencies

Conditional probabilities

Types of stimulus-response relations

THE TEMPORAL STRUCTURE OF ELICITED BEHAVIOR

Elicitation and response sequences

Adjunctive behavior

Interim and terminal behavior

Effects of successive elicitations

Adaptation or habituation

Facilitation or potentiation

Effects of time since the last eliciting stimulus

The role of exercise

The walking reflex

SUMMARY AND EXTENSION

Stimulus presentations in imprinting

Motivation and the significance of stimuli

THREE

When a stimulus elicits or reliably produces a response, this stimulus-response relation is called a reflex. The Indo-European root *steig-*, to stick, is an ancestor of *stimulus*. The same root also generated *distinguish*, *instinct*, and, via *stylus*, a writing instrument, *style*. Through the Old French *estiquet*, to impale and later to label, it also produced a modern French word now related to behavior, *etiquette*. The Indo-European root *spend-*, to pour a libation or to make a treaty, led to the Greek *sponde*, a drink-offering, and the Latin *spondere*, to promise. Through these words, *response*, originally an undertaking in return, is linked to *spouse*, *sponsor*, and perhaps even to *spontaneous*. The latter relation is interesting because *response* now refers to a unit of behavior that need not be produced by a stimulus; responses may be elicited by stimuli but they may also occur spontaneously, when they are said to be emitted.

This chapter deals with the effects of presenting stimuli. First we will examine some relatively simple situations in which a stimulus produces or elicits a response. Then we will consider how elicited behavior is structured in the time between stimulus presentations; the topics of adjunctive, interim and terminal behavior are concerned with the order in which responses occur when two or more different responses are generated by a stimulus. The next topics, adaptation and potentiation, deal with the ways in which elicited behavior may change over successive stimulus presentations. Repeated presentations of an eliciting stimulus may also lead to behavior that no longer depends on stimulus presentation; this possibility is explored in a section on the law of exercise. The chapter closes with a discussion of ways in which the behavioral significance of a stimulus may change over time; imprinting is an example used to relate such effects to the concept of motivation.

THE REFLEX

A simple way to change an organism's behavior is to present a stimulus. For example, if someone is engaged in a quiet conversation, a sudden loud noise will probably stop the talking and produce the change in posture called the startle

reaction. This reliable relation between an environmental event, a stimulus, and a resulting change in behavior, a response, has been called a *reflex*. The vocabulary of the reflex has a long history (Fearing, 1930), travelling through a complex route from the philosopher Descartes to its contemporary usage. Descartes, a seventeenth-century philosopher, was familiar with hydraulic devices constructed to amuse visitors in the royal gardens of France. Stepping on a concealed trigger released a flow of water that made statues move. Descartes saw a similarity between such devices and the reflex. According to Descartes, stimuli were comparable to

strangers, who entering into one of the grottoes containing many fountains, themselves cause, without knowing it, the movements which they witness. For in entering they necessarily tread on certain tiles or plates, which are so disposed that if they approach a bathing Diana, they cause her to hide in the rosebushes, and if they try to follow her, they cause a Neptune to come forward to meet them threatening them with his trident. (Descartes, translated in Fearing, 1930, pp. 20-21)

Just as a step on the concealed plate triggers the movement of a statue, a stimulus triggers a response. For Descartes, the role of pipes and water in these statuary systems was played by nerves and animal spirits in living organisms.

For our purposes, the most important part of Descartes' concept of the reflex was that the term incorporated the observation that behavior is sometimes caused by environmental events. Thus, a child who touches a flame will quickly withdraw its hand. The development of physiology as a science turned attention toward the mechanism of such behavioral relations, and experiments explored the components of the reflex arc, the path from the original sensory impact of the stimulus through the central nervous system and then back to the muscular or glandular system within which the response occurred. The analysis of the reflex became more and more sophisticated (e.g., Sherrington, 1906), and in the conditioned-reflex concepts of Pavlov (1927) and the related behaviorism of Watson (1919), it became the basic behavioral unit (see also Phillips, 1971).

We noted two illustrations of reflexes above, in the startle reaction and in withdrawal from a flame. Many others are familiar: the knee jerk produced by a tap on the patellar tendon; salivation produced by food in the mouth; pupillary constriction produced by bright light; the postural adjustments produced by an abrupt loss of support. These examples have the common feature that some stimulus reliably produces some response, and this is the defining property of a reflex: *the reliable production of a particular response by a specified stimulus*. In these circumstances, we say that the stimulus *elicits* the response, or that the response is *elicited* by the stimulus; the stimulus is an *eliciting* stimulus, and the response is an *elicited* response.

The relation between the two events defines the reflex (Skinner, 1931). The reflex is not a stimulus, nor is it a response. For example, we could not speak of a reflex if we flashed a bright light but observed no pupillary constriction. The light alone is insufficient (we would not be tempted to discuss reflexes if the light were flashed with no organism present). Neither could we speak of a reflex if we

observed a response without an identifiable stimulus. By itself, the response of pupillary constriction is not a reflex. We should not confuse pupillary constrictions produced by bright light with those produced in other ways. The pupillary reflex is defined by the relation between bright light and pupillary constriction, and reflexes in general are defined by the relation between eliciting stimuli and elicited responses. Many reflexes have been named for their characteristic responses; the salivary reflex and the patellar reflex are examples. But it is useful to remember that these are *not* names for responses; if we observed salivation or knee jerks in the absence of eliciting stimuli, the reflex terminology would be inappropriate.

PROPERTIES OF ELICITED BEHAVIOR

Once we identify a reflex relation between some stimulus and some response, we can examine properties of the reflex (cf. the Laws of the Reflex; Skinner, 1938). Consider an eliciting stimulus such as an acid or sour-tasting solution on the tongue (e.g., vinegar) and an elicited response such as salivation. Above some minimum value, called the *threshold*, acid on the tongue will elicit salivation, but a very low concentration or a very small quantity may not do so; when the magnitude of a stimulus is not sufficient to elicit a response, the stimulus is said to be below threshold. As we concentrate the solution or increase its quantity, or both, we will reach values that are more and more likely to produce salivation. At a given concentration, we can find out how many drops must be placed on the tongue to elicit salivation reliably; conversely, for a given number of drops, we can find out what concentration is required. Some concentrations and quantities will have only marginal effects, sometimes eliciting salivation and sometimes not.

Note that a stimulus does not itself have a threshold. Rather, we determine threshold values for particular *features* of stimuli. When we do so, we hold other features of the stimulus constant. For example, in eliciting salivation, the threshold number of drops would typically be smaller for a strong acid solution than for a weak acid solution; it would therefore not be meaningful to specify the threshold number of drops without also specifying the acid concentration, or vice versa.

Elicited responses have varied properties. Some time must elapse between stimulus and response; this time period is called the *latency* of the response. In addition, the response must exist in some *magnitude* and have some *duration*. Because these properties may vary together, they have sometimes been given a common name, reflex *strength*. Thus, reflex strength would be weak if responding occurred with long latency, small magnitude and short duration, whereas it would be strong if responding occurred with short latency, large magnitude and long duration.

The eliciting effect of a stimulus does not involve any direct transfer of energy from the environment to the organism. Sechenov (1863), a Russian physiologist, noted that the organism's energy expenditure in many responses (e.g., sneezes and coughs) far exceeds the energy provided by the eliciting stimulus. Thus, observed Sechenov, the stimulus should be regarded as a kind of trigger,

releasing energy that the organism already has available in muscles or glands or other structures. This view was, of course, consistent with Descartes's conception of the reflex.

Even though the eliciting stimulus is most accurately regarded as a trigger that releases energy already available in the organism, the magnitude of the eliciting stimulus may affect the elicited response. Response latency typically varies *inversely* with stimulus magnitude; in other words, response latency decreases as stimulus magnitude increases. And response magnitude and response duration typically vary *directly* with stimulus magnitude; in other words, these measures increase as stimulus magnitude increases. Saying that reflex strength increases with stimulus magnitude simply summarizes these relations.

Once it became common to speak of reflex relations in terms of strength rather than in terms of specific measures, it also became easier to regard the reflex relation as a fundamental unit of behavior: reflexes began to be treated as building blocks from which more complex behavior was constructed. The concept of the reflex had a tempting simplicity, and as it came to be more widely accepted as a behavioral unit, it seemed reasonable to conclude that reflex relations might be a basis for understanding a variety of behavioral processes. Some stimuli had been identified as causes of some responses, and the faulty assumption was made that for *every* response there must exist a corresponding eliciting stimulus.

Pavlov's conditioned-reflex system and Watson's behaviorism of the 1920s and 1930s grew out of such an assumption. With the reflex serving as a unit of behavior, complex behavior was then treated as the combination of reflex units. When a response occurred with no observed eliciting stimulus, stimuli with appropriate properties were hypothesized. In addition, the responses of one reflex were assumed to have stimulus properties that enabled them to elicit other responses in turn. Thus, behavior that extended over long periods of time could be interpreted as a sequence or chain of reflexes, in which each response functioned simultaneously as the elicited response of one reflex and the eliciting stimulus of the next. These reflex systems were elaborated in various ways, but despite the ingenuity and the dedication of their proponents, they no longer command substantial attention in the psychology of learning. The concept of the reflex has a place in the analysis of behavior, but by itself it is not adequate for thoroughly describing what organisms do.

Eliciting stimuli and response probabilities

Of the possible relations between stimuli and responses, the reflex is one specialized relation among many. In a reflex, the presentation of some stimulus reliably produces some response. But the stimulus of that reflex may affect other responses differently, and the response of that reflex may be affected differently by other stimuli. For any particular stimulus, its presentation may increase the likelihood of some responses, decrease the likelihood of others, and have no effect on still others. For any particular response, its likelihood may be increased by some stimuli, be decreased by others, and be unaffected by still others (Catania, 1973b).

Furthermore, it is not enough simply to note that a response reliably follows a stimulus to conclude that the language of the reflex is appropriate. We must also know how likely the response is without the stimulus. For example, if a rat in a running wheel spent much of its time running and it ran in noise as well as quiet, we would not want to say that noise had elicited running just because we observed the rat running when we presented noise. To specify the relation between a response and a stimulus we must know how likely the response is in the absence of the stimulus as well as in its presence.

Consider watching a dog. The dog may occasionally move about, prick up its ears or bark. If we give the dog some food, those responses may cease as eating begins. If we then introduce a strange dog, the eating will stop and the first dog may growl and assume an aggressive posture. If the other dog is removed and the first dog's master enters, we may observe barking and jumping and tail-wagging. But if the master now scolds, the dog may whine and its tail may drop down between its legs. The list of possible effects of various stimuli could easily be extended, but these examples are sufficient. The food, the strange dog, the master and the master's scolding each make some responses more likely while they make others less likely. Some of the responses might be observed from time to time even in the absence of these stimuli, and none will necessarily occur every time a particular stimulus is presented. In a reflex relation, a response that is infrequent in the absence of some stimulus occurs reliably when that stimulus is presented. As the above examples indicate, this class of stimulus-response relations is only one case among a variety of possible stimulus-response relations. The response of barking, for example, is affected by several different stimuli, and we may be interested in stimuli that make it less likely as well as those that make it more likely.

Probabilities or relative frequencies. These various relations among stimuli and responses can be described best in quantitative terms. We must introduce the concept of probability, or relative frequency. We will then be able to define the effects of stimuli on the basis of both the probability of a response when a stimulus is present and the probability of that response when the stimulus is absent. A probability or relative frequency is a proportion or ratio: the number of times an event occurs compared with the number of times it could have occurred. For example, lightly stroking the bottom of a newborn infant's foot typically elicits a splaying or spreading of the toes; this relation is called the Babinski reflex. We may estimate response probability by counting how many times a response is produced over a given number of stimulus presentations. In the newborn infant, the proportion of the toes may be produced by each of twenty touches to the foot; the proportion of responses is therefore 20 out of 20, or a probability of 1.0. The Babinski reflex ordinarily diminishes with age, and if we test this reflex again at a later date, only 6 out of 20 touches may now produce a response; the probability is 0.3. Still later in the infant's life, we will probably find that the reflex has completely disappeared: none of the 20 touches produces a response, and the probability is therefore 0.0. In the mathematical notation for probabilities, these examples may be

written respectively as:

$$p(R_1) = 1.0; p(R_2) = 0.3; \text{ and, } p(R_3) = 0.0.$$

Here, p stands for probability, and the terms in the parentheses are abbreviations for the events whose probabilities are specified in the equation (in this instance, R for responses, where the subscripts 1, 2, and 3 correspond to the three successive measurements of the Babinski reflex). These examples also illustrate that probabilities are defined in such a way that they can have values only within the range from 0.0, when the event never occurs, to 1.0, when the event always occurs.

Conditional probabilities. The Babinski reflex is an uncomplicated example for our purposes, because the toe-splaying response does not occur often in the absence of its eliciting stimulus. Suppose, however, that we study the blinking elicited by a puff of air to the eye. If we study this reflex with an infant whose blinking already occurs at the rate of a blink every one or two seconds, how can we tell elicited blinks from those that would have occurred even in the absence of the air-puff? In fact, we may never be able to say whether or not any particular blink is elicited. But by comparing the probability of a blink after an air-puff with the probability after no air-puff, we may at least be able to assess the effect of the stimulus. Our procedure is illustrated in Figure 3-1. We watch the eyelid and record blinks within a period of time after each stimulus, and we compare responding in these time periods with responding in equivalent time periods that do not follow a stimulus.




Figure 3-1. Estimating eye-blink probability with and without an eliciting puff of air. Each solid vertical line represents a blink. In Line A, no air-puffs were presented, and the dashed lines mark off five one-second periods during which blinks were recorded. Blinks occurred within three of these five periods, and the probability of a blink in the absence of a stimulus is therefore 0.6. In Line B, each arrow represents a puff of air to the eye. An eye-blink occurred in each of the five one-second periods that followed these stimuli. The probability of a blink given an air puff is therefore 1.0. The two blinks produced by stimulus X, in time-period 3, did not both count toward this probability; a time period is scored either as one in which no responding occurred or as one in which at least a single response occurred, and probability is then calculated by dividing the number of time periods with at least one response by the total time periods. These data can be summarized by saying that the air-puff raised the probability of an eye-blink from 0.6 to 1.0. (This example is hypothetical; an actual experiment would use a much larger sample of observations to estimate probabilities.)

The probabilities in this example are expressed in the following notation:

$$p(R/S) = 1.0; p(R/\bar{S}) = 0.6.$$

In this notation, R is a response, S is stimulus, and a bar over a symbol (as in $\bar{S}$) is a negative sign, equivalent to *no* or *not*. Thus, the first equation may be read as: probability of a response given that a stimulus has been presented (eye-blink given air-puff) is 1.0. The second equation may be read as: probability of a response given that a stimulus has not been presented (eye-blink given no air-puff) is 0.6. This kind of probability, in which the probability of one event is specified in terms of the presence or absence of another event, is called a *conditional probability* (the terminology should not be confused with that of Pavlov's conditioned or conditional reflexes, even though both cases have the common feature that one event is a condition for some other event). Within the parentheses, the term to the left of the slash represents the event whose probability is measured, and the term to the right represents some other event upon which the first is conditional. In other words, $p(A/B)$ can be read as: probability of A given B ; or, probability of A on the condition that B has occurred. Similarly, $p(X/\bar{Y})$ can be read as: probability of X given not- Y ; or, probability of X on the condition that Y has not occurred.

Later discussions will favor the language of probabilities over other ways of describing behavior, but usage in psychology varies enough that such choices of vocabulary will often be arbitrary. There is not much difference between saying a stimulus produced a response and saying a stimulus elicited a response. These usages would be applied most often to single instances (e.g., this particular presentation of stimulus M was followed by this particular instance of response N). Neither is there much difference between saying a stimulus increased response likelihood and saying a stimulus raised response probability. These usages would most often be applied to the average effect of a stimulus over many instances (e.g., stimulus O is usually followed by response P).

Types of stimulus-response relations

Now we may summarize some relations between stimuli and responses in terms of conditional probabilities. For this purpose, we will use the coordinate system in Figure 3-2. The y -axis is a scale of response probability given that stimulus has been presented; the x -axis is a scale of response probability given that the stimulus has not been presented.¹ In other words, any point on this graph represents two conditional probabilities: response probability given a stimulus, $p(R/S)$, and response probability given no stimulus, $p(R/\bar{S})$. For example, the point labeled A in the left graph of Figure 3-2 represents a stimulus-response relation in which response probability is 0.75 when the stimulus is presented and only 0.50

¹The reader may elsewhere have encountered the terms *ordinate* and *abscissa* in place of the terms y -axis and x -axis. Strictly, the terms are not interchangeable. Whereas x -axis and y -axis refer to the dimensions of the coordinate system, *abscissa* and *ordinate* refer respectively to particular locations on the x -axis and on the y -axis. Thus, the response of the above example, with a probability of 0.75 in the presence of a stimulus and a probability of 0.5 in its absence, is represented in Figure 3-2 by the intersection of an ordinate of 0.75 and an abscissa of 0.5.




Figure 3-2. Various stimulus-response relations represented in terms of response probability in the presence of a stimulus, $p(R/S)$, and in the absence of the stimulus, $p(R/\bar{S})$. The left graph shows examples of stimulus-response relations in which a stimulus raises response probability (A), has no effect on response probability (B), or lowers response probability (C); the increase in response probability called a reflex, in which the stimulus reliably produces the response, is illustrated at D. The right graph shows possible stimulus-response relations that might be observed with various combinations of responses (a dog's salivation, growling or tail-wagging) and stimuli (its master, food or a strange dog). This graph makes the point that stimuli can have different effects on different responses, and that responses can be differently affected by different stimuli.

when the stimulus is absent; in this instance, the stimulus raises the probability of the response.

The diagonal in these graphs is of particular interest. A response that occurs without being elicited by a stimulus is said to be *emitted*. The diagonal represents stimulus-response relations in which response probability is independent of or is unaffected by the stimulus. Thus, at point B response probability is 0.25 whether or not the stimulus is presented. In such cases, it would be inappropriate to say that a stimulus produced or elicited the response. To say a response has been elicited it is not enough to know that the response followed a stimulus. A response can happen to be emitted just after a stimulus is presented. Thus, when a response follows a stimulus, we cannot say that the response was elicited unless we know that the stimulus actually produced or caused the response; otherwise, we must say that the response was emitted.

A third class of stimulus-response relations is illustrated by point C, for which the response probability of 0.90 when the stimulus is absent is reduced to 0.10 by the stimulus; in this instance, the stimulus makes response probability decrease. (In some behavioral vocabularies, such reductions of response probability by a stimulus are called *reflex inhibition*; for example, if a dog's tail-wagging stopped with the

introduction of a strange dog, we might say that the strange dog inhibited tail-wagging.) All points above and to the left of the diagonal represent cases in which a stimulus increases response probability, whereas those below and to the right represent cases in which a stimulus decreases response probability.

Among those cases in which a stimulus raises response probability are those in which virtually every stimulus presentation produces the response. For example, at point *D* a stimulus raises response probability from about 0.1 to about 1.0. This is the kind of stimulus-response relation we have been calling a reflex. The graph demonstrates that this relation is only one extreme in a range of possible stimulus-response relations. Somewhere between the upper edge of the graph, where $p(R/S)$ is near 1.0, and the diagonal, where *S* has no effect on *R*, we might decide that the eliciting effect of the stimulus is no longer reliable enough to justify calling the stimulus-response relation a reflex. But where should the boundary be? We probably would not exclude cases in which $p(R/S)$ is just a little less than 1.0 (e.g., 0.98, or perhaps even 0.95). But neither would we want to include very slight effects of stimuli as reflexes. For example, would we wish to speak of a stimulus-response relation as a reflex if a response with a probability of 0.2 were raised by a stimulus only to 0.21, or even to 0.25?

Any boundary we could draw to separate those stimulus-response relations we call reflexes from those we do not will be arbitrary. This is an important conclusion. We saw how some earlier behavioral systems, such as those of Watson and Pavlov, based their development on the reflex as a fundamental unit of behavior. No doubt highly reliable reflex relations were easier to work with than other, less reliable stimulus-response relations. But if Figure 3-2 is appropriate for describing stimulus-response relations and if the reflex is only a special case among these relations, then any system of behavior built solely on the reflex as a behavioral unit was bound to leave something out.

One property of behavior left out of early accounts was the emission of responses, the occurrence of responses in the absence of eliciting stimuli. Emitted responses were given such names as *instrumental* or *operant* behavior, because they were studied in terms of how they were instrumental in producing changes in the environment or how they operated on the environment. They derived their importance not from their relation to eliciting stimuli but rather from their consequences. By contrast elicited behavior was called *reflex* or *respondent*. Once this distinction had been made, other features were added. In particular, it was argued that instrumental or operant behavior consisted of skeletal responses, such as movements of the limbs, whereas reflex or respondent behavior consisted of autonomic responses, such as glandular secretions. This distinction was also seen as paralleling the traditional distinction between voluntary and involuntary action. Only recently have these distinctions been seriously questioned. Skeletal responses can be elicited and glandular responses can be emitted. It is important to maintain the distinction between elicited and emitted responding. But we have seen that the same response may be sometimes elicited and sometimes emitted; thus we cannot classify responding effectively into these two categories on the basis of such

physiological properties of behavior as the difference between skeletal and autonomic responses.

Now let us return to Figure 3-2, to consider specific cases of stimulus-response relations. The right graph of Figure 3-2 shows hypothetical relations involving three stimuli and their effects on three responses. A dog's salivation and growling and tail-wagging are represented by three locations on the x-axis, and the master, food or a strange other dog are represented by filled circles, unfilled circles, and unfilled triangles, respectively. The probability of salivation is raised both by the master and by food but is lowered by the strange dog; these stimuli have similar effects on tail-wagging, except that food elicits salivation more effectively than does the master whereas the master elicits tail-wagging more effectively than does food. Conversely, growling becomes more probable in the presence of the strange dog and less probable in the presence of the master or food; we might also say that food has a larger inhibitory effect than the master.

This graph supports an earlier point. When we defined the reflex, we argued that the reflex is neither stimulus nor response but rather the relation between them. The graph represents stimulus-response relations; it cannot represent stimuli by themselves or responses by themselves. Thus, food as a stimulus does not have any single location in the graph; its location depends on the response we are measuring. Similarly, tail-wagging as a response does not have any single location in the graph; its location depends on the stimulus we are presenting.

THE TEMPORAL STRUCTURE OF ELICITED BEHAVIOR

We have seen some complexities of analyzing the relations between stimuli and responses. Of the possible relations between stimuli and responses, the reflex is one specialized relation among many. In a reflex, a specified stimulus reliably produces a particular response. But the stimulus of that reflex may affect other responses differently, and the response of that reflex may be affected differently by other stimuli. For any stimulus, its presentation may increase the likelihood of some responses, decrease the likelihood of others, and have no effect on still others. For any response, its likelihood may be increased by some stimuli, decreased by others, and unaffected by still others. We now move on to examine how elicited responding changes over time.

Elicitation and response sequences

Presenting a stimulus may determine not only whether a response occurs at a particular moment but also the sequence of responses that occurs over an extended time. For example, if a rat is deprived of food and we give the rat a food pellet, the rat will eat the pellet. Once it has finished eating the pellet, the rat will then typically drink water if water is available. This relation between eating and subsequent drinking is so strong that by delivering food in small amounts over an extended period we can make the rat drink several times its ordinary daily ration of water (Falk, 1961, 1971).

Adjunctive behavior. Falk named this increase in drinking *polydipsia*, and he called responding that depended in this way on other responding *adjunctive* behavior. Adjunctive behavior, therefore, is behavior in which one response reliably accompanies some other response produced by a stimulus. Polydipsia, as one type of adjunctive behavior, follows at least partly from the way in which eating and drinking are related in the rat's normal feeding pattern. If food and water are freely available, the rat ordinarily takes a few large meals daily and drinks after each meal. If we then force the rat to take many small meals by delivering food in many small portions at intervals of one or a few minutes, the rat still drinks after each meal but does not reduce the size of each drink enough to compensate for its more frequent drinking. Thus, a rat that went from five large meals to fifty small meals per day would now drink ten times as often. But if drinks taken after each of the fifty small meals were only half instead of one-tenth the size of drinks taken after each of five large meals, the rat would drink five times as much as before. Such increases in the rat's daily water intake can be readily produced, and polydipsia is so reliable in rats that it has been used to get rats to ingest substances that they ordinarily reject (e.g., alcohol: Meisch and Thompson, 1971).

Adjunctive behavior includes other responses besides eating and drinking. For example, if a running wheel is available to a rat instead of water, running in the wheel will follow eating in much the same way as drinking follows eating in polydipsia (Levitsky and Collier, 1968). Another example involves aversive stimuli. If electric shock is delivered to a monkey's tail, the monkey will typically bite anything nearby it can sink its teeth into. The electric shock, in other words, elicits biting. Once the biting stops, the monkey then will typically engage in manual responses, such as grabbing and shaking levers, chains or other manipulable objects (Hutchinson, Renfrew, and Young, 1971). The manual responses may be regarded as adjunctive behavior relative to the shock-elicited biting.

Interim and terminal behavior. The presentation of stimuli may impose temporal order or structure on behavior. A stimulus may generate a sequence of different responses. Figure 3-3 illustrates an experiment that examined patterns of behavior generated in pigeons by repeated food presentations (Staddon and Simmelhag, 1971). Each pigeon was placed in a chamber on one wall of which was an opening to a feeder. The feeder consisted of a tray of grain ordinarily out of the pigeon's reach; food was presented by lighting the tray and lifting it to a position where the pigeon could eat from it. Figure 3-3 illustrates a pigeon's behavior during constant 12-second periods that separated successive 2-second food deliveries.

The experimenters observed the pigeon through a window in the chamber and scored responding in various categories. Some categories were movements, such as pecks directed toward the feeder wall or toward the floor; others were postures that could be maintained for some time, such as orientation toward the feeder wall. The categories were determined by watching many pigeons for an extended period.

Figure 3-3 shows that early in each 12-second interfood period, Bird 29 usually made quarter-turns, and then, but less frequently, put its head into the feeder opening or pecked toward the floor. Other pigeons showed patterns that included different responses, such as preening or pacing. These responses became




Figure 3-3. Probability of a pigeon's responses in successive seconds of a 12-second period between successive food deliveries. Food deliveries consisted of 2-second access to grain in a food hopper located in one of the chamber walls. Data are based on the last three of 109 sessions of 64 food deliveries each. Orientation toward the feeder wall, not shown, increased in probability over each 12-second period. Low-frequency responses such as wing-flapping and preening are also omitted. A sequential response pattern is implied by the different times at which each response reaches its maximum probability. (Adapted from Staddon and Simmelhag, 1971, Figure 1)

less likely and pecking became the dominant or most likely response as the time of the next feeder presentation approached. Staddon and Simmelhag called the varied responding early in each period *interim* behavior (in Figure 3-3, quarter-turns, head-in-feeder and floor-directed pecks). The more consistent responding later in each period, which typically included some form of pecking, was called *terminal* behavior (in Figure 3-3, wall-directed pecks).

The original purpose of Staddon and Simmelhag's experiment was to evaluate an earlier experiment on repeated food presentations by Skinner (1948). Skinner observed that pigeons tended to repeat those responses that had just preceded earlier food presentations, and referred to this repetition as an experimental *superstition*.

The conditioning process is usually obvious. The bird happens to be executing some response as the hopper appears; as a result it tends to repeat this response. If the interval before the next presentation is not so great that extinction takes place, a second "contingency" is probable. This strengthens the response still further. . . . The experiment might be said to demonstrate a sort of superstition. The bird behaves as if there were a causal relation between its behavior and the presentation of food, although such a relation is lacking. (Skinner, 1948, pp. 168-171)

Skinner noted that the pigeon's responding often changed gradually as the procedure continued. He referred to such gradual changes in responding as *topographical drift* because they depended on accidental relations that developed between responding and food deliveries; the procedure did not seem to force changes in the topography or form of responding in any particular direction. Although a pigeon might respond temporarily as if its responses were producing food as a consequence, its responding would change or drift over the long run, and no particular response could be identified consistently as superstitious behavior. Over longer periods of observation, Staddon and Simmelhag observed that pecking eventually dominates as the response just preceding food deliveries in this kind of procedure.

Superstition remains somewhat controversial, and we will return to it later. For the moment, our concern is with the temporal structure that repeated stimulus presentations impose on behavior. Responding early in interfood intervals, interim behavior, varied from one pigeon to another, whereas later responding, terminal behavior, was fairly constant across pigeons and usually included some form of pecking. Pecking is also the behavior occasioned by food presentations; the bird takes food into its mouth by pecking. At least in this case, terminal behavior seems to have something in common with responses produced by stimulus presentations. (Similarly, if a monkey wears a collar with a leash, electric shocks will elicit biting and pulling at the leash, but repeated shock presentations make the biting and pulling occur as terminal responses, in the time just preceding each shock: Morse, Mead, and Kelleher, 1967). It would be nice to conclude that repeated stimulus presentations not only elicit responses but also produce terminal behavior closely related to the elicited responding. Unfortunately, behavior is not that simple. Researchers have not yet settled on how or even on whether the classes of behavior here called adjunctive or interim and terminal or superstitious should be distinguished. The problem is not about the facts of behavior, but about how these facts should be organized.

Effects of successive elicitations

We have described some effects of stimulus presentations on some properties of behavior. One complication in our description is that one presentation of a given stimulus may not have the same effect as a later presentation of the same stimulus. For example, we may startle much more to the first flash of lightning in a thunder

storm than to later flashes of lightning. Furthermore, the effects of stimuli may depend on how quickly they follow each other. For example, if we are peeling onions, the tears elicited by the present onion may depend on whether we began working on it immediately after finishing the last one or had just returned after taking a five-minute break. In other words, to describe how stimuli can change behavior, it will not be enough merely to list the changes in responding produced by a stimulus. Elicited responding may depend on the number of stimulus presentations or their spacing in time. This may be the point at which elicitation first becomes relevant to learning. If we show only that some stimulus elicits some response, we have not demonstrated learning. But we may wish to begin speaking of learning if the responding elicited at one time depends on what happened to the organism earlier.

Adaptation or habituation. We already mentioned the startle reaction produced by an unexpected event such as a lightning flash or a sudden loud noise. Even without other events that signal its occurrence, a repeated loud noise will produce successively smaller startle reactions, until eventually no behavior at all is observable when the noise occurs. Many stimuli elicit responses called orienting or observing responses; for example, a dog pricks up its ears in response to a novel sound or begins sniffing in response to an unusual odor. With repetitions of these stimuli, however, the dog's responding decreases; the responding occurs with smaller magnitude and longer latency, until perhaps it vanishes completely.

This decrement in responding with repeated stimuli has been given various names. For the stimulus-response relations called reflexes, this decrement has been called *adaptation* or *habituation*. The terms have complicated histories, and the distinctions between them remain ambiguous. Both physiological and behavioral considerations enter into their usages, and the different terms are invoked on the basis of such circumstances as whether the observed decrements in responding can be attributed to specific kinds of changes in the nervous system (cf. Groves and Thompson, 1970). For our purposes, the terms *adaptation* and *habituation* are interchangeable.

Adaptation or habituation is a characteristic of the elicited responding produced by a variety of stimuli. It has been studied with such diverse responses as the change in skin resistance produced by electric shock (the galvanic skin response or GSR), distress calls of birds to the silhouette of a predator passing overhead, and contractions in earthworms produced by exposure to light, not to mention the startle reactions and orienting responses mentioned earlier (e.g., Ratner, 1970). It may also be an important component of the dynamics of emotion (R. L. Solomon and Corbit, 1974).

Facilitation or potentiation. But stimuli may sometimes have opposite effects. For example, electric shocks elicit squealing in rats; if several shocks are delivered, later presentations produce more responding than earlier ones (e.g., Badia, Suter, and Lewis, 1966). This effect has been called *facilitation* or *potentiation*. Either is an acceptable term for increased responding with repeated presenta-

tions of a single stimulus.² The process of facilitation or potentiation seems more characteristic of stimuli regarded as aversive or punishing than of stimuli regarded either as neutral or as appetitive or reinforcing.

The method of stimulus presentation can determine whether adaptation (habituation) or facilitation (potentiation) is observed. Kimble and Ray (1965) studied a wiping reflex in the frog: a bristle touched to a frog's back elicits a movement over the back by the hind foot. In one group of frogs, successive touches were made within a specified region of the frog's back each time, but within this region the exact location of stimulation could vary from one touch to the next; in a second group, successive touches were made to precisely the same location on the frog's back each time. In both groups, 100 touches at 10-sec intervals were delivered each day for 12 consecutive days. The two groups began with roughly equal probabilities of elicited wiping movements, but this probability increased over successive days for the first group whereas it decreased for the second group. In other words, facilitation or potentiation occurred when the location of the eliciting touch varied slightly from presentation to presentation, whereas adaptation or habituation occurred when the location of the eliciting touch was constant.

In another example, Hutchinson, Renfrew, and Young (1971) examined biting responses elicited in squirrel monkeys by electric shock. When shocks were delivered at 30-sec intervals, the elicited biting responses gradually decreased over a 60-min session. When shocks were delivered at 120-sec intervals, these responses remained roughly constant over the session. And when shocks were delivered at 240-sec intervals, responding gradually increased over the session. Thus, the same stimulus, depending on its frequency of presentation, could elicit responding that demonstrated adaptation, facilitation, or neither of these phenomena.

Effects of time since the last eliciting stimulus. Repeated presentations of a stimulus change the extent to which that stimulus elicits responses. If the stimulus is then absent for some time, the probability that responding will be elicited may return to earlier values. For example, the startle reaction to loud noise may diminish or even disappear after several loud noises in succession over a period of a few minutes, but may appear again in full strength if a loud noise later follows hours of silence. Such an effect of the time between stimuli follows from the nature of adaptation and facilitation. If elicited responding decreased over successive stimuli, we would expect it to recover to its earlier higher levels after the stimuli were discontinued. Conversely, if elicited responding increased over successive stimuli, we would expect it to return to its earlier lower levels after the stimuli were discontinued. The changes in responding called adaptation and facilitation are not permanent, and the return to earlier levels takes place as time passes. In fact,

² Another term also used occasionally is *sensitization*, but sensitization has an alternative and more widely accepted usage. It refers to an increase in the eliciting effect of one stimulus as a result of presentations of some other stimulus; one stimulus amplifies the eliciting effect of another stimulus (for example, a flash of light may make it more likely that a later loud noise will produce a startle reaction).




Figure 3-4. Temporal effects of presenting stimuli. Arrows indicate stimuli; vertical lines indicate probabilities with which these stimuli elicit responses. Effects usually called adaptation or facilitation are illustrated in *A*, *B* and *C*; responding decreases in probability with successive stimuli. Effects sometimes called facilitation or potentiation are illustrated in *D*, *E* and *F*; responding increases in probability with successive stimuli. In both cases, response probabilities return to earlier levels as the time since the last stimulus increases (*A* through *C* and *D* through *F*).

this temporary nature of the decrease or increase in response probability is part of the definition of adaptation and facilitation.

These relations are summarized in Figure 3-4. Each line shows hypothetical effects on response probability or reflex strength of ten successive stimuli. The upper part of the figure (*A*, *B* and *C*) illustrates adaptation or habituation; the

lower part (*D*, *E* and *F*) illustrates facilitation or potentiation. In both cases, three examples are given that differ only in the time separating the first five stimuli from the second five stimuli. As this time becomes longer (from *A* to *C* and from *D* to *F*), the responding produced by the second five stimuli becomes more like the responding produced by the first five stimuli. In other words, responding returns to its earlier levels with the passage of time. For example, in *A* responding is elicited by the later stimulus *x* with a lower probability than by the first stimulus *w*. But when stimulus *x* is presented after successively longer periods without stimuli, as in *B* and *C*, the probability with which stimulus *x* elicits responding approaches that with which stimulus *w* elicited responding. Analogous relations exist for stimuli *y* and *z* in *D*, *E* and *F*.

The role of exercise

We argued above that adaptation and facilitation are not ordinarily permanent. If decreases in responding were irreversible, we either could observe adaptation only once in an organism's lifetime, or else successive adaptations would drive responding to lower and lower levels until it finally disappeared altogether. A parallel case for facilitation would similarly lead either to a single instance of facilitation in the organism's lifetime or to continuing and unlimited increases in elicited responding. None of these alternatives seems plausible.

Yet consider the following experiment by Dill (1974) with a tropical fish, the zebra danio. Dill measured the zebra danio's escape response from a stimulus designed to resemble the rapid approach of a larger predatory fish. As this stimulus approached, the danio began to swim away from it (despite its aquatic medium, this approach, the danio began to swim away from it (despite its aquatic medium, this swimming response of the fleeing fish could aptly be called a flight reaction). The latency with which swimming began (or the threshold for size or rate of growth of the stimulus) decreased with repeated experiences; this is equivalent to saying that the stimulus more and more strongly elicited the escape response. But when tested after 10 days without further elicitations, the response had not returned to its earlier values. The change produced by the stimulus was not reversible. Facilitation or potentiation are therefore not appropriate names for the effect. How then shall we treat it, and what are its implications?

"Practice makes perfect" is a familiar saying about the role of repetition in behavior. Before the effects of the consequences of responding were appreciated, it was believed that the mere repetition of responding, without regard to its consequences, was sufficient to maintain behavior. Consider the following from Sechenov:

... an infant is able to cough, sneeze and swallow immediately upon birth. The act of sucking also belongs to this category of complex movements. . . . Moreover, Indeed, everybody knows that a new-born child is able to suck. . . . More-
over, it is a well-known fact that the activity of this complex mechanism in the infant is called forth by irritation of the lips; put, for example, your finger, or a candle, or a wooden stick between the child's lips, and it will begin to suck. Try to do the same with a child three months after it has been weaned,

and it will no longer do so; however, the ability to produce sucking movements at will is retained by man for life. These are highly remarkable facts: on the one hand, they show that the conduction of sensation from the lips to the central nervous mechanisms which produce the sucking movements apparently ceases in the child after weaning; on the other hand, they indicate that the integrity of this conduction is maintained by the frequent repetition of the reflex. (Sechenov, 1863, pp. 28-29)

Sechenov here emphasizes the complexity of the response in the sucking reflex, but more important is his observation on the role of repetition. Not only is repetition important in maintaining the response, according to Sechenov, but the response also becomes independent of the effects of eliciting stimuli. In the infant, sucking is elicited by stimuli ("irritation of the lips"); later, these stimuli no longer elicit the response, but the organism through adulthood remains able to produce the response even in the absence of these stimuli (cf. Schoenfeld, 1966, Segal, 1972). In our vocabulary of stimulus-response relations, we might say that the repeated elicitation of a response increases the likelihood that the response will be emitted.

Early accounts of learning (e.g., Thorndike, 1913; cf. Verhave, 1967) treated effects of response repetition as basic components of learning, described in terms of laws of *exercise* or laws of *practice*. These laws were often ambiguous on such questions as whether it mattered if the repeated response were elicited or emitted. In any case, they soon were overshadowed by other concerns. As the psychology of learning turned to other phenomena, such as reinforcement and stimulus control, the possibility of a role for exercise or practice became neglected. The evidence is too scanty to allow firm conclusions about whether exercise or practice might be a basic component of learning. Nevertheless, some tantalizing data exist, and some speculation is therefore in order.

Some of the earliest responses in an organism's lifetime occur spontaneously, in the absence of identifiable eliciting stimuli. For example, inside its egg the embryo chick makes uncoordinated movements of its limbs and body. These movements may prevent the developing bones from becoming fixed in their sockets, or may modify the form of the growing bones and connective tissue. Later in the embryo's development, eliciting effects of stimuli appear, perhaps simply as the embryo's sensory apparatus matures. This progression from spontaneous responding to elicited responding may be summarized by saying of chick embryos that "they 'act' before they 'react.'" (Provine, 1976, p. 210). Stimuli become more important later, as when the chick's rotating movements in pecking its way out of its shell during hatching are affected by whether the chick continues to encounter intact portions of the shell.

Evidence suggests that pecking in the young chick depends not only on the conditions that elicited pecks and the consequences of earlier pecks but also on how much pecking the chick has already engaged in (Hogan, 1971). Once responding occurs, whatever its origin, it may have consequences, and the chick's survival may depend on whether those consequences in turn affect its behavior. Among gulls, for example, pecking at but missing the parent's beak has different consequences from striking the beak; only in the latter case is the parent gull likely to feed the

chick. These differential consequences will affect the accuracy of the chick's later pecking only as the response becomes independent of its eliciting stimuli. That development corresponds to one possible formulation of a law of exercise: once a response has been elicited by a stimulus, the response may become more likely even in the absence of the stimulus³ (cf. Spalding, 1873, on the relation between instinct and learning).

In experiments on salivation, dogs at first salivate only when food is presented, but after several food presentations they begin also to salivate occasionally when food is absent (e.g., Zener and McCurdy, 1939). Such responding, called *spontaneous* salivation, has been attributed to conditioning of the salivary response to features of the experimental setting. For example, it might be argued that the eliciting stimulus was a spot on the wall that the dog just happened to see when food deliveries began. No other eliciting stimuli are identifiable, so the only way to salvage the treatment of salivary responses as part of a reflex relation is to assume that salivation could be elicited by such arbitrary features of the experimental setting. Yet the problem of identifying a stimulus simply vanishes once we admit the possibility that responses can occur without eliciting stimuli; in fact, we cannot otherwise conceive of emitted responding at all. The concept of emitted responding will be essential to our treatment of consequential operations such as reinforcement and punishment.

Another example of a transition from elicited to emitted responding comes from signalled avoidance in rats. In these experiments, a warning stimulus such as a buzzer precedes electric shock. If the rat responds during the buzzer and before shock, the shock is omitted; if the rat responds after shock has begun, the shock is removed. Thus, during the warning stimulus the rat can *avoid* shock by responding. *ing*, whereas during the shock the rat can *escape from* shock by responding. A frequent experimental choice of response in rats is some form of locomotion, such as jumping a hurdle or running from one side of the chamber to the other; the avoidance response is typically the same as the escape response. But shock often elicits responses like these in rats, and once such responses have been produced by shock, they may also occur when shock is absent. Thus, a rat's first few avoidance responses may occur because the elicitation of such responses by shock makes responses likely to be emitted in the presence of the warning stimulus (cf. Azrin, Hutchinson, and Hake, 1967; Hutchinson, Renfrew, and Young, 1971, Figure 12). Once avoidance responding begins, it may continue because it has the important consequence of preventing shock.

The walking reflex. Early behavior theories held that the simple repetition of behavior was important in its own right to the development and maintenance of that behavior. Consider the following from Sechenov: "... If a child which has just learned to walk becomes ill and remains in bed for a long time, it forgets the previously acquired art of walking. . . . This fact testifies once again to the great import for nervous activity of frequent repetition" (Sechenov, 1863, p. 29). We can

³ A change in the probability of a response in the absence of a stimulus with repeated elicitations of the response implies a shift in the point at which the stimulus-response relation is located in the coordinates of Figure 3-2; beyond recognition of this implication, our current understanding of these phenomena does not justify further speculation.

no longer be sure of the evidence upon which Sechenov based his conclusion (e.g., did he observe one child or many children?). But we can consider a contemporary example, in Zelazo, Zelazo and Kolb (1972a, b).

These investigators worked with the walking reflex, which appears in newborn infants and disappears at about eight weeks. The response of the walking reflex has much in common with the behavior that the child engages in later when learning to walk. Coordinated walking movements resembling those of an adult can be elicited by holding the infant under the arms and letting its feet touch a level surface. In one of four groups of six male infants each, parents exercised their infant's walking reflex by holding him so as to elicit the walking response; four brief sessions of exercise were held daily, for a total of ten minutes per day (active-exercise group). In a second group, parents exercised their infant's limbs for the same time per day, by moving the arms and legs with a gentle pumping motion while he was lying on his back; the walking response, however, was not elicited (passive-exercise group). The infants in the two remaining groups were not exercised systematically. These served as control groups, to assess effects of weekly testing of the walking reflex and of the regular periods of attention given to the infants during exercise.

On the average for all infants, the walking response was elicited at a rate of five to ten responses per minute during the first week of life. For all but the active-exercise group, this responding dropped to a rate of less than four per minute by the eighth week; for the active-exercise group, elicitation of the walking response increased gradually over the eight weeks to almost thirty responses per minute. Furthermore, the mean age at which walking began was 10.1 months for infants in the active-exercise group, 11.4 months for those in the passive-exercise group, and 11.7 and 12.4 months for those in the two control groups. Thus, exercise of the walking reflex during the infant's first eight weeks of life not only increased the frequency of an elicited response that ordinarily decreases in frequency during this time period; it may also have shortened the time to the later appearance of this response as a component of emitted behavior, walking.

The active-exercise and passive-exercise groups controlled for effects of parental attention, but Zelazo, Zelazo and Kolb recognized that walking movements might produce other changes in the infant's world (visual, kinesthetic, or tactile) that could reinforce such responses. They also noted that even during the eight weeks of exercise the walking response was sometimes emitted rather than elicited, and they took this as evidence for learning:

Walking among the active-exercise infants seemed to progress from a reflexive to an instrumental response. There is little doubt that learning occurred. The behavior of one 5-week-old infant, in particular, is difficult to dispute; he made vigorous walking motions while the observer suspended him in midair over the table. . . . Not only were there more responses by the active-exercise infants, but they were better executed. (Zelazo, Zelazo, and Kolb, 1972b, p. 315)

We have speculated that after a response is produced by a stimulus it may become more likely even in the absence of the stimulus. In other words, eliciting a

response may raise its probability of emission. This formulation differs from the classical laws of exercise or practice to which we have related it. With the evidence at hand, we cannot say how it will stand up to future research. Nevertheless, such a process is undoubtedly the simplest that could enter into instances of learning, because it depends only on the simplest of experimental operations: the presentation of stimuli.

SUMMARY AND EXTENSION

Up to this point, we have concentrated on how stimuli affect responding. Virtually all of the phenomena considered so far can be defined and described in terms of response probabilities: in a reflex relation, a stimulus raises the probability of a response to near 1.0; in reflex inhibition, a stimulus reduces the probability of a response; in adjunctive behavior, the elicitation of one response changes the probability of some other response; in adaptation, response probability decreases over successive stimulus presentations; in facilitation, response probability increases over successive stimulus presentations; and in the possibly controversial phenomenon of exercise, successive elicitations raise response probability in the absence of the stimulus. We now examine an outcome of stimulus presentations, *imprinting*, that must be discussed in terms other than the effects of the stimulus on response probability (cf. Lorenz, 1937; Hess, 1973). Imprinting includes some effects of stimulus presentation that provide a bridge to the treatment of response consequences in the next chapter.

Stimulus presentations in imprinting

When a duckling hatches, the first moving thing it is likely to see is its mother, and even on this first day of its life outside the egg the duckling will probably begin to stay close to her. But if the mother is gone and the duckling first sees something else in motion, such as a human, the duckling will behave toward this stimulus as it otherwise would have toward its mother. Such stimuli are said to be *imprinted*, i.e., in a figurative sense, stamped into the duckling.

Imprinting has been demonstrated in both laboratory and field with a variety of stimuli, ranging from real and model birds to electric trains (some stimuli, of course, work better than others). The development of imprinting is sometimes said to have a *critical period* of one or a few days: if imprinting does not occur during this critical period, it may not occur at all.⁴ Whatever the basis for the critical period, the duckling begins to respond in significant ways to a stimulus, whether mother duck, human, or some arbitrary moving object, if the stimulus is introduced early enough in the duckling's life. One of these responses is following the

⁴One view (Hoffman and Ratner, 1973) suggests that this happens because fear of novel stimuli develops by the end of the critical period; as the birds get older, they make characteristic distress calls in the presence of and move away from novel stimuli. Thus, older birds do not ordinarily stay near such stimuli long enough for imprinting to occur. This view implies that imprinting could occur at any time in the duck's life if this later effect of novel stimuli could be prevented or reversed.

imprinted stimulus as it moves; it has sometimes been said that this following is elicited by the imprinted stimulus, but this usage of the language of elicitation is misleading.

We might say that the imprinted stimulus, through its presentation to the duckling in appropriate circumstances, has acquired control over the duckling's behavior. For example, if the imprinted stimulus is the mother duck, the duckling will follow her about and will emit distress calls in her absence. But what kind of control does the imprinted stimulus exert over the following response? Whenever the duckling walks toward the mother, it finds itself closer to her; whenever the duckling walks from the mother, it finds itself farther away. In other words, the natural consequence of walking in different directions is to change the distance between the duckling and its mother. If closeness to the mother is an important consequence for the duckling, it should be no surprise that it walks toward her rather than away from her. It also follows that if the duckling's environment is changed so that the presence of the mother requires some response other than walking, that other response should begin to occur in its place.

Such an experiment was designed by N. Peterson (1960). On one side of a window was a dark compartment containing a moving imprinted stimulus. A duckling, on the other side of the window, could light up the side with the imprinted stimulus by making an appropriate response. It did so even when the response was one incompatible with following, such as pecking at a disk located on the wall or standing still on a platform near the window. In other words, the critical property of the imprinted stimulus was not that it could elicit particular responses, such as following or pecking or standing still, but rather that it had become an important consequence for the duckling and therefore could reinforce or increase the frequency of such responses as following or pecking or standing still. In the natural environment, following will ordinarily keep the duckling close to the imprinted stimulus, but a laboratory environment shows that the duckling will learn to make other responses if they are the ones having this important consequence.

In imprinting, the effects of stimulus presentation are not changes in the probabilities of responding; rather they are changes in the significance of the stimulus that has been presented. The imprinted stimulus acquires its significance for an organism simply by its presentation at a particular time in the organism's life. It begins as a stimulus toward which the organism is relatively indifferent but ends as an important stimulus, one that will affect the duckling's responding when it is made a consequence of behavior.


Motivation and the significance of stimuli

There are other ways to change the significance of stimuli. For example, if we knew that a rat was more likely to eat or drink than to run in a running wheel, we might expect that the rat would press a lever more often if its lever presses produced food or water than if its presses produced only access to the running wheel. While continuing to allow the rat free access to food and water, however, we could remove the wheel, or lock it so that it could no longer rotate, thereby

preventing the rat from running in the wheel. After a period of deprivation of wheel running, we might then find that the rat would be more likely to run than to eat or to drink if the opportunity to run in the wheel was available again, and that the rat would press the lever more often if presses produced access to the running wheel than if they produced food or water. In other words, by depriving the rat of one or another of these stimuli, we can change their significance relative to one another: we can make running more likely than eating, or drinking more likely than running, or eating more likely than drinking, depending on the rat's prior deprivation of each.

Changes in the significance of stimuli are the basic concern of the study of *motivation* (e.g., Bolles, 1975). Stimuli may be made more or less reinforcing or more or less aversive, depending upon such factors as the time since their last presentation. Food, for example, becomes more reinforcing as time passes since the last opportunity to eat, and it may even become aversive after eating if an unusually large quantity is consumed. (The significance of stimuli can also be changed through means other than deprivation; for example, conditioned reinforcers and conditioned aversive stimuli are stimuli that have acquired their reinforcing or aversive properties through their relation to other stimuli. For a discussion of the interaction between motivation and adaptation effects, see R. L. Solomon and Corbit, 1974.)

We have reviewed many aspects of the stimulus-presentation operation, and are ready to move on to the consequential operations of reinforcement and punishment. The effects of stimuli may be important to our understanding of consequences. Whenever responses have consequences, these consequences will also have their stimulus effects. With all of the changes in responding that can be produced by presenting stimuli, we should not expect to understand consequences without taking them into account.



Consequences of Responding: Reinforcement

THE LAW OF EFFECT

Mazes and learning curves

Experimental chambers and cumulative records

REINFORCEMENT

The relativity of reinforcement

Extinction

Extinction versus inhibition

Contingencies and stimulus deliveries

Extinction and superstition

THE ACQUISITION OF BEHAVIOR

Latent learning

Sensory-motor learning

SUMMARY

FOUR

Through the Indo-European root, *sekw-*, to follow, *consequence* is linked to *sign* and *designate*, from the Latin *signum*, something that one follows, and to *social* and *association*, from the Latin *socius*, a companion or follower. It shares its prefix, *con-*, with *conditioning*, *contingency*, and *contiguity*. *Conditioning*, through the Indo-European root *deik-*, to show or pronounce, has many relatives: *teach*, from the Old English *tæcan*, to show or instruct; *judgment*, from the Latin *judex*, he who pronounces the law; and *paradigm*, from the Greek *para*, beside, and *deiknūnai*, to show. *Contingency*, from the Latin *contingere*, to touch on all sides, to happen, has varied meanings: a possibility; a condition of being dependent on chance; something incidental to something else. Like contact, it combines the roots *com-*, together, and *tangere*, to touch. *Contiguity*, the condition of touching or being in contact, has the same origins. Curiously, *contingency* and *contiguity* are usually contrasted in psychology: *contingency*, in its technical use, stresses how the likelihood of one event may be affected or caused by other events, whereas *contiguity* simply implies the juxtaposition of events in space or time, without regard to causation.

The environment can do things to an organism, but the organism in turn can do things to its environment. An organism's behavior can have consequences. A rat without food may find food by going to some other location. Merely by moving from one place to another, the rat changes its environment from one without food to one with food. But food as a consequence may also change the rat's future behavior: the rat may go there again later when deprived of food. It is an important fact of behavior that it can be affected by its consequences.

Arranging consequences for behavior is not simply presenting stimuli. Rather, the stimuli must be presented in some relation to behavior. The environment must be arranged so that the organism's responses will be followed by specified events. Consequences for behavior already exist in the natural environment. Even before we intervene, organisms can change their environments by doing things or by going from one place to another. But we can better study how consequences affect behavior by arranging consequential operations in the laboratory. We can build a maze so that a water-deprived rat will find water after making an appropriate sequence of turns. We can build a chamber within which a food-deprived pigeon can produce food by pecking a key on the wall. Then we can see how water affects the turns the rat takes as it runs through the maze or how food affects the rate at

which the pigeon pecks the key. In both cases, responses have consequences and the consequences may affect future responding.

This chapter begins by exploring the historical development of experiments on the consequences of behavior. It then treats some properties of the contemporary concept of reinforcement: the relativity of reinforcement, the significance of discontinuing reinforcement (extinction), the range of consequences that can be reinforcing, and the range of responses that can be reinforced. It concludes by showing that reinforcement is relevant to simple sensory-motor interactions like the visual consequences of eye movements as well as to behavior maintained by physiologically significant consequences such as food or water.

THE LAW OF EFFECT

Chapter 2 considered Thorndike's experiments with animals in problem boxes, from which the animals could escape by operating an appropriate device that released the door. Typically, the food-deprived animal was placed inside the box with food available outside. In its varied activity, the animal sooner or later operated the device that released the door and was then free to leave the box. With repetition the animal operated the device more and more rapidly upon being placed in the box. The response, at first of low probability, became more and more probable as it repeatedly had the consequence of opening the door. Thorndike described how the consequences of responding affected later responding in terms of a principle he called the *Law of Effect*. The law went through many revisions, but its essence was that responses could be made more probable by some consequences and less probable by others. In a language closer to Thorndike's, responses with satisfying effects would be stamped in whereas those with annoying effects would be stamped out.¹

Figure 4-1 illustrates data from one of Thorndike's cats. To escape from this box, the cat had to pull a string that ran from a wire loop at the front of the box to a bolt that held the door. The first time in the box, the cat took 160 seconds to escape. This time decreased gradually and irregularly over successive trials until the cat was reliably escaping in less than 10 seconds during the last few trials. This process, the gradual decrease in the time taken to complete a task, came to be called trial-and-error learning (Köhler contrasted this gradual change with the sudden or insightful solutions he observed with chimpanzees).

Mazes and learning curves

In later years, trial-and-error learning was studied with many different organisms in many different kinds of situations. Experimenters believed that the intelligence of different species could be compared by how rapidly learning

¹This version was called the *strong* Law of Effect. Later, Thorndike repudiated the second half of the Law, leaving the increase in probability or stamping in but discarding the decrease in probability or stamping out. What remained was then called the *weak* Law of Effect. This historical point will be relevant in chapter 5, when we deal with punishment.

proceeded in puzzle-boxes, mazes, runways, and other diverse apparatus (e.g., Hilgard, 1951; Woodworth and Schlosberg, 1954). Apparatus design began to be dictated by theoretical questions: whether learning took place in discrete steps, on an all-or-none basis, or instead occurred gradually and continuously; whether organisms learned movements (response learning) or properties of the environment (stimulus learning); whether the consequences of responding led directly to learning, or only made the organism perform so as to demonstrate what it had learned in other ways. Some of these issues were the basis for long-lasting controversies (Goldstein, Krantz, and Rains, 1965).

A common feature of these various experiments was that responding could be made more probable when it had certain consequences. The change in probability was measured differently depending on the apparatus and the experimental aims. Graphs showing how behavior changed in the course of an experiment were called *learning curves*: the time to complete a response as a function of number of trials (e.g., Figure 4-1); the percentage of correct responses; the proportion of animals reaching some criterion of successful performance at different stages. Sometimes these measures were transformed to ease comparisons among them. For example, in successive runs of several rats through a maze, each rat's time to run from startbox to goalbox decreases while each rat's percentage of correct turns and the proportion of rats making errorless runs increases. Converting the time to run through the maze to speed, defined as the reciprocal of the running time, made it possible to compare three different measures all of which increase with learning. But the shapes of learning curves depended so much on the apparatuses used and




Figure 4-1. A learning curve. A cat's time to escape from a problem box as a function of trials. (From Thorndike, 1898, Figure 1)

the measures taken that the progress of learning could not be described in any unitary way.

The problem was that the performances in these experiments were complicated. For example, measuring the time course over which a rat stopped entering blind alleys as it learned its way through a maze did not show how learning proceeded at a single choice point. This consideration led to the gradual simplification of mazes, as illustrated in Figure 4-2. Diagram *A* shows the plan of the earliest maze used for the study of animal learning (Small, 1899-1900). The maze was a six-by-eight-foot modification of the hedge maze at Hampton Court in England. (Curiously, such mazes may also have provided the setting for the hydraulically operated statues that contributed to Descartes's concept of the reflex.) The rat's cage was placed at the start, and when the cage door lifted the rat could enter the maze; food was in the goal area at the center. With increasing experience in this maze rats reached the goal area more rapidly and with fewer wrong turns along the way. But it was difficult to examine learning at any particular choice point. If the choice at 1 was learned more quickly than the choice at 7, the difference might be due either to the early and late positions of these choice points in the sequence or to the different floor-plans of these choices; if the choices at 4 or 5 were learned more quickly than the choice at 3, the difference might be due to the two different ways in which 4 could be approached (from 3 or from 5) or to the fewer occasions on which 5 would be encountered if the rat often went from 3 to 4 directly rather than by way of 5.

Gradually, maze structures evolved into more systematic forms, as in *B* of Figure 4-2. In this maze, sometimes called a *U*-maze after the form of the successive units, choice points were essentially the same as the rat approached each of them; they differed only in their locations in the sequence and in whether left turns or right turns were correct. This kind of systematic arrangement made it easy to specify the correct sequence (in *B*, right-left-right-left-left-right) and to keep track of errors. Even here, however, interactions among positions and sequences complicate the analysis. For example, is the rat's choice of left at 4 affected by whether it is preceded by a right turn at 3 and is followed by another left turn at 5; would it matter if the rat approached 4 after it came back from the blind alley at 3, having made an error there, rather than after a correct right turn at 3; is it important that 4 is in the middle of the maze rather than near its beginning or its end?

It was perhaps inevitable that the maze would be reduced to a single choice point, as in the *T*-maze with a goalbox on the right at *C* in Figure 4-2. Here, when the rat left the startbox it had only to make a single choice of right or left. But possible complications were not yet eliminated. For example, suppose that one rat in its first trial in the *T*-maze turned right whereas a second rat turned left. Should the second rat be allowed to retrace its steps once it reaches the empty box at the end of the left arm? If instead it is returned to the startbox, should it be forced to the goalbox, by blocking the left arm of the alley, to make sure its experience in the goalbox is comparable to the first rat's? The next logical step was to eliminate the choice points completely, leaving nothing but a simple runway, as at *D* in




Figure 4-2. Stages in the evolution of mazes in the study of animal learning: A. the Hampton Court maze, as adapted by Small (1899-1900); B. a U-maze with six choice-points; C. the single-choice-point T-maze; and D. the runway or straight alley.

Figure 4-2. Now no errors were even possible, and the measures of behavior were reduced simply to the speed with which the rat moved from startbox to goalbox.

There were other problems. An average measure of the performance of a group of animals did not necessarily represent the performance of individuals in the group. Suppose, for example, that for single rats in a simple T-maze frequent errors typically change abruptly after several trials to consistently correct turns, but this change occurs on different trials for different rats. In a large group of rats, 65% might make correct turns by the fifth trial, and 72% by the sixth, and 79% by the seventh, and 83% by the eighth, and so on, until performance became stable at 98% to 100% by the twentieth trial. This group performance would give the appearance of a gradual increase in correct turns, and would completely obscure the abrupt change in performance in the behavior of individual rats (cf. Sidman, 1952; Estes, 1956).

Even the simple runway was not an ultimate solution, because the speed of running down a straight alley could be affected by many trivial factors. If trials began with the opening of a startbox door, speed might be affected by the direction the animal was facing when the door opened. Speed might also be affected by the

experimenter's handling of the animal when moving it between trials from the goalbox to the startbox, or by odor trails left by other animals, or even by the room in the goalbox if it forced the animal to slow down so as not to bang its head against the goalbox wall.

With mazes or runways, the experimenter had to return the organism from goalbox to startbox to begin a new trial, and was only indirectly recording behavior when measuring time taken to get from startbox to goalbox. In such cases, the experimenter rather than the organism determined when behavior could occur, and the overall time measure was not equivalent to specifying in detail the particular sequence of responses on a given trial. Two experimental innovations helped to solve these problems. First, an apparatus was designed in which the organism could repeatedly emit an easily specified response without the experimenter's intervention. Second, responding was measured directly, by its rate or frequency, rather than indirectly by other measures derived either from complex response sequences or from groups of organisms. Inspired partly by an interest in reducing the handling of the organism and thereby simplifying the work of the experimenter, these were important features of a direction of research initiated by Skinner (1930, 1938, 1950; see esp. Skinner, 1956, for a history of these developments).

Experimental chambers and cumulative records

Two representative apparatuses are illustrated in Figure 4-3: a standard rat chamber with a single lever on the left and a three-key pigeon chamber on the right. They share response devices (the rat's lever and the pigeon's three keys), mechanisms for delivering reinforcers such as food or water (a pellet dispenser for the rat and a grain hopper for the pigeon), and stimulus sources (a speaker and a lamp for the rat and lamps or projectors behind each of the keys for the pigeon; the rat chamber also often includes a grid floor for the delivery of electric shock).

In a typical arrangement with a rat, the rat is placed in the chamber when food-deprived. A lever the rat can press protrudes from one wall. Near the lever is a food cup into which food pellets can be dispensed from a delivery system on the other side of the wall; the operation of the pellet feeder usually produces a distinctive sound, and the pellet delivery may be accompanied by the brief lighting of the feeder cup by a small lamp (not shown in Figure 4-3). General illumination is provided by the houselight, and noise can be broadcast from the speaker to mask sounds from outside the chamber.

The first step is feeder training. Pellets are delivered into the food cup, and first it is necessary to wait until the rat finds and eats the pellets. Once this happens, pellet deliveries are continued until upon each delivery the rat comes quickly to the food cup from any chamber location; ten or twenty pellet deliveries are ordinarily sufficient. Once feeder training is completed, the apparatus is changed so that pellet deliveries depend upon lever presses. Sooner or later the rat presses the lever, this lever press produces a food pellet, and the food pellet occasions eating. The rat will then probably go back to the lever and press it again. (Alternatives to waiting for the lever press are considered in chapter 6.) The




Figure 4-3. A rat chamber (left) and a three-key pigeon chamber (right). The rat chamber includes a lever (A), a food cup and pellet delivery tube (B), a speaker (C), a lamp or houselight (D), and a grid floor through which electric shock can be delivered (E). The pigeon chamber includes three keys (F, G, H) and the opening to a food hopper (I). Lamps or projectors behind each key allow colors or patterns to be displayed on the keys.

outcome of interest is the frequency with which the rat presses the lever. If this frequency increases, we call the food pellet a *reinforcer*. In the type of chamber shown in Figure 4-3, other kinds of reinforcers could be substituted. For example, the pellet dispenser might be replaced by a dipper that delivers small quantities of water or milk.

A pigeon chamber differs from a rat chamber in that keys are substituted for levers and the feeder accommodates the pigeon's standard diet. The key is simply a piece of plastic mounted behind a round hole in the chamber wall. One component of the key is a switch that allows electrical recording of the pigeon's pecks if they are forceful enough (the keys are routinely sensitive to forces of less than 0.1 Newton, which is about 10 grams or one-third of an ounce; for pigeons' key pecking, this sensitivity is adequate for most purposes). The plastic is usually translucent, so that lamps or miniature projectors behind the key can project patterns or colors on it. The pigeon chamber in Figure 4-3 contains three keys, arranged horizontally about 23 centimeters (9 inches) above the chamber floor. Any particular experiment might use only a single one of these keys, some combination of two, or all three. Typically, keys are lit when they are in use.

The opening to the feeder is centered below the keys. The feeder consists of a tray of pigeon food (mixed grain or commercially available pigeon pellets) that can be brought within the pigeon's reach when operated. Like the delivery of the rat's pellet, the operation of the pigeon feeder is accompanied by characteristic sounds; it is also common practice to light the feeder and turn off all other chamber lights whenever the feeder is operated. The chamber typically includes other features, such as a houselight for dim general illumination, sources of masking noise and of other auditory stimuli, ventilation and temperature control, and so on.

As with the rat, if a food-deprived pigeon's pecks on a key produce food, we may expect an increase in the rate at which the pigeon pecks that key. In these

apparatuses, the opportunity to eat is used to raise the probability of a response. Whether the response is a rat's lever press or a pigeon's key peck, the important common property of these arrangements is that responses can have consequences and that these consequences can change behavior. The rat and the pigeon are common laboratory organisms. They each have idiosyncratic species-specific behavior patterns that must be taken into account, and we must not take it for granted that what we observe with rats or pigeons can be generalized to other organisms. Nevertheless, their diet, housing, susceptibility to disease, and other characteristics are reasonably well understood, and their small size, relatively long life-span, and economy make them particularly convenient. Thus, we will find that they have often served in experiments on the consequences of responding.

The responding in apparatuses like those in Figure 4-3 has sometimes been called *free-operant* responding: *free* because the organism is free to emit the response at any time rather than waiting for the experimenter (as when the rat in a maze goalbox must wait to be placed back in the startbox before it can run through the maze again); and *operant* because the response operates on the environment. Free-operant responding lends itself to a recording method, the cumulative record, that provides a detailed and convenient picture of how responding changed over time. In a cumulative recorder, illustrated in Figure 4-4, a roll of paper is




Figure 4-4. Principal components of a cumulative recorder. The roller drives the paper at a constant speed, and each response moves the pen a fixed distance across the paper. The paper speed and the step-size for each response vary with the behavior under study. A common scale is about 1 centimeter per minute (about 2.5 minutes per inch) and 1100 responses across the width of the paper (about 200 responses per inch). At this scale, a slope of 45° represents a rate of about 40 responses per minute. When the pen moves to the top of the paper, it resets automatically to its starting position near the bottom.




Figure 4-5. Some representative cumulative records. Records *A* and *B* differ mainly in rate of responding: rate is higher in *A* than in *B*. The rate is zero throughout most of record *C*; a segment of *C* during which a few responses occurred is magnified and shown in relation to an event record in *D*. Records *E* and *F* are roughly equal in response rate, but show different detailed patterns of responding: *E* is steplike, indicating periods of responding that alternate with pauses, whereas the smoother-grained *F* indicates relatively steady responding. Records *G* and *H* show rates that change over time: response rate decreases in *G* (negative acceleration) whereas it increases in *H* (positive acceleration). In this figure, a slope of 45° represents a rate of about 20 responses per minute.

threaded around a roller. A motor drives the roller at a constant speed, feeding out the paper. A pen or other writing device rests on the paper as it passes over the roller, and each response (e.g., a rat's lever press or a pigeon's key peck) moves the pen a small distance along the roller, at right angles to the movement of the paper. The resulting record therefore shows the total responses accumulated at any time during the session.

Examples of cumulative records are shown in Figure 4-5. Because the paper moves at a constant speed, the higher the rate of responding the steeper is the slope of the cumulative record, as illustrated in records *A* and *B*. In the scale for Figure 4-5, the rate of responding is roughly 30 responses per minute for record *A* and roughly 12 per minute for record *B*. Record *C* includes only a few responses; the horizontal portions indicate periods of time when no responses occurred (note that a cumulative record cannot have a negative slope, because the pen can record responses only by moving in one direction across the page.)

A small section of record *C* that includes a few responses is magnified in record *D*, in which the magnified cumulative record is accompanied by an event record on the same time-scale. For each response in the event record, a small step occurs in the cumulative record; this property of cumulative records is not usually obvious because typical response and time scales are too small for such fine resolution of detail. Nevertheless, different patterns of responding can be distinguished in cumulative records. For example, the rates of responding in records *E* and *F* are roughly the same but record *E* is steplike whereas record *F* is relatively smooth. This means that the responding that produced record *E* occurred in short high-rate bursts (steep segments) separated by pauses (flat segments), whereas the responding that produced record *F* occurred at a much more uniform rate from moment to moment. This property of cumulative records is sometimes called *grain*; of the two records, *E* has the rougher grain and *F* has the smoother one.

Records *G* and *H* provide other examples of the detailed properties of behavior made visible in cumulative records. In record *G*, the rate begins at roughly 25 responses per minute, but as time passes it gradually decreases; in record *H*, the rate changes in the opposite direction, increasing from a relatively low rate to roughly 30 responses per minute (such patterns of responding are respectively called negatively and positively accelerated, because the slopes of the records respectively decrease and increase over time).

Figure 4-6 shows additional features sometimes incorporated into cumulative records. Records *A* and *B* show how displacements of the response pen can be used to indicate other events besides responses. In these records, only some responses produced food, irregularly in *A* (as at *a*, *b* and *c*), and regularly in *B* (as at *d* and *e*). The repeated concave pattern in *B*, as illustrated between *d* and *e*, is sometimes called scalloping. In record *C*, responding that began at *f* and continued through *g* was followed by some event (e.g., food delivery) indicated by the pen displacement at *g*. The pen then reset to *h* and the sequence repeated, as at *h* to *i* and so on. This type of record makes successive segments easy to compare (e.g., more responding occurred in the segment ending at *g* than in the one ending at *i*). Record *D* illustrates how sustained pen displacements can distinguish different conditions. For example, in the presence of a tone the pen may be in its normal position, as in segments *j*, *l*, and *n*, whereas in the absence of the tone, the pen may be displaced downward, as in segments *k*, *m*, and *o*.

With this treatment of free-operant behavior and cumulative records, we have explored part of the technological development of the science of behavior. Before we can effectively consider the findings made available through such analyses, we must turn our attention to some aspects of the language of behavior.

REINFORCEMENT

When a food-deprived rat's lever presses produce food, lever pressing becomes more probable than when the presses do not produce food. When a water-deprived pigeon's key pecks produce water, key pecking becomes more probable than when the pecks do not produce water. And perhaps when a child's crying produces the




Figure 4-6. Additional features of cumulative records. In *A* and *B*, displacements of the pen are used to superimpose a record of other events, such as food deliveries, on the record of cumulative responses (as at *a* through *e*). In *C*, the resetting of the pen is used to simplify comparisons among successive segments of the record (*f* to *g*, *h* to *i*). In *D*, sustained displacements of the pen are used to distinguish responding in the presence of one stimulus (at *j*, *l*, and *n*) from nonresponding in the presence of another stimulus (at *k*, *m*, and *o*). In this figure, a slope of 45° represents a rate of about 40 responses per minute (note that this scale differs from that in Figure 4-5).

parent's attention, crying becomes more probable than when it does not produce attention. These cases illustrate the principle of reinforcement: when responding produces reinforcers, responding increases. The principle is simple, but as it evolved from Thorndike's initial versions of the Law of Effect to its contemporary status it has carried problems of language and logic with it (Catania, 1969b, 1973b). Some properties of the contemporary vocabulary of reinforcement are summarized in Table 4-1.

The vocabulary of reinforcement includes the term *reinforcer* as stimulus and the terms *reinforce* and *reinforcement* as either operation or process. For example, when the rat's lever presses produce food pellets and lever pressing increases, we may say that the pellets are reinforcers and that the lever presses are reinforced with pellets. Although a reinforcer is a kind of stimulus, reinforcement is neither a stimulus nor response. Instead, reinforcement as an operation is presenting a reinforcer when a response occurs. This operation is carried out on responses, and therefore we speak of reinforcing responses rather than of reinforcing organisms. We may say that a food pellet reinforced a rat's lever press or that a pigeon's key

TABLE 4-1

The Vocabulary of Reinforcement.¹ This vocabulary is appropriate if and only if three conditions exist: (1) a response produces some consequence; (2) the response occurs more often than when it does not produce that consequence; and (3) the increased responding occurs *because* the response has that consequence.

<i>Term</i>	<i>Restrictions</i>	<i>Examples</i>
reinforcer (noun)	A stimulus.	Food pellets were used as reinforcers for the rat's lever presses.
reinforcing (adjective)	A property of a stimulus.	The reinforcing stimulus was produced more often than the other, nonreinforcing stimuli.
reinforcement (noun)	As an operation, the delivery of a consequence when a response occurs.	The fixed-ratio schedule of reinforcement arranged food deliveries after every tenth key peck.
to reinforce (verb)	As a process, the increase in responding that results from the reinforcement operation.	The experiment with monkeys demonstrated reinforcement produced by social consequences.
	As an operation, to deliver a consequence when a response occurs; responses are reinforced and not organisms.	When a period of free play was used to reinforce the child's completion of class assignments, the child's grades improved.
	As a process, to increase responding through the reinforcement operation.	The experiment was designed to find out whether gold stars would reinforce cooperative play among first-graders.

¹ A parallel vocabulary is appropriate to punishment (including *punisher* as a stimulus and *punish* as a verb; cf. chapter 5), with the difference that a punishing consequence makes responding occur less rather than more often.

peck was reinforced with water, but it would be incorrect to say that food reinforced the rat or that the pigeon was reinforced for pecking or that a child was reinforced. The most important reason for this restriction is illustrated in the last three examples: when we speak of reinforcing organisms, it is all too easy to omit the response or the reinforcer or both. The restriction therefore forces us to be explicit about what is reinforced by what. Nor are we forced to omit the organism; we can always say whose response it is (e.g., the child's crying).

The term *reinforcement* has also functioned often as a name for the process that follows from the reinforcement operation, i.e., the increase in responding. This dual usage, as both operation and process, complicates the description of behavior but remains established in the psychological literature. For example, the statement

that a response was reinforced can be interpreted in two ways: either the response produced a reinforcer (operation) or responding increased because it produced a reinforcer (process). This text favors the usage of reinforcement as an operation. The process that follows from this operation can be described so easily in terms of changes in responding (e.g., response probability increased) that there seems little justification for substituting other terms for a direct description of what happens to responding. Nevertheless, the process usage has so much precedent that we will be unable to avoid it completely. In addition, the overlap between operation and process vocabularies extends to a variety of other terms in learning (e.g., extinction, punishment, conditioning; cf. Ferster and Skinner, 1957; Catania, 1968).

The vocabulary of reinforcement leads to some logical difficulties even when restricted to experimental operations. When a response becomes more probable because it has produced a stimulus, we say the response has been reinforced and we call the stimulus a reinforcer. If asked how we know that the stimulus was a reinforcer, we can point to the increase in responding. If then asked why the increase occurred, we may say it did so because the response was reinforced. At some point, we begin to repeat ourselves. Once we define a reinforcer by its effect on behavior, we create a problem of circular definition if we simultaneously define the effect by the reinforcing stimulus (cf. Meehl, 1950).

One solution is to recognize that the term *reinforcement* is descriptive rather than explanatory. It names a particular relation between responses and the environment, but it does not explain that relation. The relation is complex and must include at least three components. First, a response must have some consequence. Second, the response must increase in probability (i.e., the response must be more probable than when it does not have this consequence). Third, the increase in probability must occur *because* the response has this consequence, and not for some other reason. For example, if we knew only that a response had become more probable, it would not be appropriate to say that the response must have been reinforced; the response might have been elicited by a stimulus. It would not even be sufficient that the response was now producing some stimulus it had not been producing before. We would still have to know whether responding increased *because* the stimulus was its consequence.

Suppose a child decides to experiment with a newborn sibling by pinching the infant whenever it cries. At first the infant might lie quietly, but if the older child is patient the infant will sooner or later cry. Then comes the pinch, but the pinch will be followed by more crying, and thus by more pinches, and thus by more crying, and so on. We assume that this procedure will inevitably attract the parents' attention and that they will intervene to rescue the infant. At that point the pinching may have consequences for the older child, but we must leave the scene before the parents arrive to get back to the language of reinforcement.

In this example, a response had a consequence and responding increased. Two of our criteria for reinforcement were therefore satisfied. Yet we would be reluctant to call the pinches reinforcers or to say that the child had reinforced the infant's crying with pinches. Our reason is that we do not believe that the increase in crying came about because of the consequential relation between crying and

pinches. We know enough about infants and crying and pinches to recognize that the pinches would have made the infant cry whether or not the pinches were produced by crying; the pinches by themselves would have made the infant cry even if they had been delivered without reference to crying. Our conclusion is that crying occurred because it is ordinarily elicited by pinches and not because it produced pinches. It was incidental that the first cry happened to lead to the first pinch; the infant cried when pinched whether or not it brought on the pinches itself. Whenever responses produce stimuli as consequences, the stimuli may have eliciting effects along with or instead of their effects as consequences of responding.

The vocabulary of reinforcement requires that the response have a consequence, that responding increase, and also that the increase occur *because* responding has its consequences and not for other reasons. Once these conditions are satisfied, we may describe the circumstances by saying that the response was reinforced and that the stimulus was a reinforcer. We may also assume that the stimulus will continue to be effective as a reinforcer in the future, and that it will reinforce other responses in other situations. Either assumption, however, may be incorrect. The effectiveness of reinforcers can change over time, and some stimuli may reinforce some responses but not others. For example, in everyday conversation saying "you're welcome" may reinforce someone else's "thank you," but money rather than "you're welcome" is more likely to reinforce other kinds of behavior, such as the services of a plumber or an electrician; by itself even "thank you" might not be effective. Despite these reservations, the reinforcers used in many standard experimental situations (e.g., food with food-deprived organisms) are effective reinforcers for a variety of responses; the experimenter who chooses a stimulus that will reinforce some responses but not others will sooner or later confront the difference and have to cope with it.

The relativity of reinforcement

The preceding discussion buttresses the logic of the vocabulary of reinforcement, but does not tell us how to identify reinforcers independently of their effects in the consequential operation of reinforcement. Without making a stimulus a consequence of responding we cannot say whether or not the stimulus will be effective as a reinforcer. Even demonstrable reinforcers such as food and water can vary in effectiveness depending on deprivation. With food and water continuously available to a rat, lever pressing might not be much affected by making food or water a consequence of lever presses. (In chapter 3, these effects were discussed in terms of *motivation*. We can now define motivation more precisely: when we study motivation we are concerned with what makes consequences more or less effective as reinforcers or as punishers.)

Reinforcers exist in great variety. Some are consumed; others are not. Some appear effective on the organism's first experience with them; others acquire their reinforcing properties during the organism's lifetime. Some are effective only if the organism makes physical contact with them; others are effective at a distance. No common physical properties allow us to identify reinforcers independently of their

effects on behavior, but we may find it useful to examine behavioral effects of reinforcers, such as the responding they occasion.

When a rat's lever press produces food, the food gives the rat an opportunity to eat. If we make both lever and food available to the rat at the same time the rat is more likely to eat than to press the lever. This kind of observation led to the hypothesis that the probability of one response will increase if that response gives the organism an opportunity to engage in another response more probable than the first (Premack, 1959, 1965, 1971). In other words, if response *A* is more probable than response *B*, an opportunity to engage in response *A* can be used to reinforce response *B*. According to this account, food is an effective reinforcer for a food-deprived rat's lever presses simply because eating is more probable than lever pressing.

Consider an experiment that demonstrated how the reinforcing effects of two stimuli could be reversed by independently varying the probabilities of the responses occasioned by these stimuli (Premack, 1962). A rat's running in a wheel was controlled by engaging or releasing a brake on the wheel, and was measured in number of revolutions. The rat's drinking from a drinking tube was controlled by moving the tube into or out of an opening in a stationary wall on one side of the wheel, and was measured with a drinkometer, an electrical system that recorded licks. As tested during brief periods when both responses were available, running became more probable than drinking when the rat's opportunity to run was restricted while water remained available, but drinking became more probable than running when the rat's opportunity to drink was restricted while the rat continued to have access to the wheel. In both cases, the opportunity to engage in the more probable response was a consequence that could be used to reinforce the less probable response. When running was more probable than drinking (after restricted access to the running wheel), licks became more likely if they released the brake on the wheel and allowed the rat to run than if they had no effect on the opportunity to run. Conversely, when drinking was more probable than running (after restricted access to the drinking tube), running became more likely if it produced the drinking tube and allowed the rat to drink than if it had no effect on the opportunity to drink.

This demonstration implies that reinforcers cannot be defined independently of the responses that they reinforce. In Premack's experiment, drinking reinforced running when drinking was more probable than running, but running reinforced drinking when the probabilities were reversed. According to this account, reinforcers are relative and not absolute. Their important properties are based on the responses for which they provide an opportunity. This relativity went unrecognized for a substantial time in the psychology of learning. Most learning experiments restricted attention to responses with relatively low probability (e.g., the rat's lever press) and to reinforcers that occasioned highly probable responses (e.g., food and eating). These cases were common and convenient but were special cases nonetheless. No one thought to ask, for example, whether by appropriate deprivation conditions a situation might be arranged in which an opportunity to press a lever could be used to reinforce eating. (The example is not so far-fetched.

The opportunity to eat a good meal can be an effective reinforcer, but consider how often children are persuaded to finish their dinners by making other activities depend on that eating.)

The relativity of reinforcement may be illustrated further by expanding the previous experiment to three responses. Let us add a pellet feeder to the activity wheel and the drinking tube and by appropriate deprivation operations make eating more probable than running and running in turn more probable than drinking. In these circumstances, running is reinforced by the opportunity to eat, but the opportunity to run reinforces drinking. In this situation, a response can simultaneously reinforce and be reinforced. By changing the relative probabilities of the three responses, we can make an opportunity to engage in any one an effective reinforcer with respect to either or both of the other two.¹

Earlier, we spoke of behavior in terms of hierarchies of responses. The ranking of responses within this hierarchy can be altered by stimuli, by the organism's opportunities to engage in different responses, and even simply by the passage of time. When we restrict an organism's opportunities to engage in some response (i.e., when we arrange deprivation), we make that response more probable or, in other words, we raise it within the hierarchy. We then may raise the probability of other responses lower in the hierarchy by using the opportunity to engage in the first response as a reinforcer.

We introduced reinforcers as kinds of stimuli, but we are now speaking of reinforcers in terms of responses. The treatment shifted because we found that an important property of a reinforcer was the responding that it occasioned. Premack discussed this as follows:

I have been accused of ignoring the stimulus; in fact, all the response probabilities we have measured have involved explicit stimulus operations, e.g., pellets, sucrose solutions, pinball machines, manipulanda, activity wheels, etc. . . . Even though it is desirable to take the position that a response is a means of calibrating the value a subject sets on a stimulus, and that the precise contribution which the response factors make to value is a matter for research, we should not swing so far in the other direction as to return to the traditional stimulus error. For example, tradition talks about the reinforcement value of food. But food has no value; we must talk about the value of food when it can be smelled (but not seen), seen (but not smelled), both seen and smelled, eaten, etc. It is part of the advantage of the response language to avoid this kind of error automatically. (Premack, 1971, pp. 125-27)

Like the reflex, reinforcement is a relation. The relation includes responding, its consequences, and the change in behavior that follows. We need not treat reinforcement as a theory or a hypothesis; it is simply a name for this relation.

¹The detailed operation of Premack's principle depends on how probabilities are calculated. Choice among simultaneously available responses seems more satisfactory than the proportion of time occupied by each response (but cf. Dunham, 1977).

Extinction

The consequences of many responses remain reasonably constant throughout an organism's lifetime. For example, a fairly reliable consequence of speaking is hearing our own voice; a fairly reliable consequence of reaching for an object nearby is touching it; and a fairly reliable consequence of walking is getting from one place to another. But for other responses, consequences change. Responses reinforced during childhood may no longer be reinforced in adulthood. Educational systems often arrange consequences such as praise or grades for spelling or completing problems in arithmetic or carrying out research projects, but sooner or later these artificial consequences are discontinued (with the hope that other more natural consequences will maintain these responses after the student moves on to other things). When a response is reinforced, it increases in probability. But the increase is not permanent, because responding decreases to its earlier levels when reinforcement is discontinued.

The operation of discontinuing reinforcement is called *extinction*, and when responding decreases to its earlier level as a result of this operation the response is said to be extinguished. This return of responding to its level before reinforcement mainly demonstrates that the effects of reinforcement are temporary. As long as reinforcement continues, responding is maintained; after reinforcement is discontinued, responding is no longer maintained. Thus, the decrease in responding during extinction is not a special process requiring a separate treatment; rather, it is part of the process generated by the reinforcement operation. The rapidity of the decrease shows how much the effects of reinforcement persist over time (Morse, 1966).

At one time, responding during extinction was a primary measure of reinforcement. *Resistance to extinction* was expressed as the time elapsed until responding dropped to some specified low rate or as the number of responses emitted during a specified period of extinction. Two hypothetical records of extinction of a rat's lever presses are shown in Figure 4-7. In both records response rate decreases as time passes (negative acceleration), but depending on the method of measurement either might represent greater resistance to extinction. If the criterion is the time until two minutes elapse without a response, then resistance to extinction for *A* is greater than for *B*; at no point does *A* include two minutes without responses whereas such a period begins halfway through *B*. If instead the criterion is the number of responses emitted, then resistance to extinction for *B* is clearly greater than for *A*. Because the definition of resistance to extinction is not unambiguous, the measure diminished in significance. Nevertheless, it remains appropriate to treat extinction in terms of the persisting effects of reinforcement; resistance to change (of which resistance to extinction is a special case) is an important property of behavior (cf. Nevin, 1974).

Extinction versus inhibition. If extinction did not occur, the effects of reinforcement would be permanent. Any response established by reinforcement would last throughout the organism's lifetime. Clearly that is not the case. For example,




Figure 4-7. Two hypothetical cumulative records of the extinction of a rat's lever presses after food reinforcement. Either *A* or *B* might be said to demonstrate the greater resistance to extinction, depending on whether it is measured in terms of the time taken until two minutes elapse without a response or in terms of the total responses emitted during the extinction session.

if we wear a watch we often turn our wrist to look at it; the consequence of looking is finding out the time. But if we stop wearing the watch for some reason, we eventually stop looking; the consequence of seeing a bare wrist is not effective as a reinforcer.

The history of the concept of extinction, however, was not so simple. It came to be argued that extinction did not merely allow the passive decline of responding but rather that it somehow actively suppressed responding. Extinction was said to have *inhibitory* effects. The view that some process actively reduces responding during extinction long dominated the alternative view that responding decreases during extinction because it is no longer maintained by reinforcement. Thus, texts on learning tended to devote separate chapters to reinforcement and extinction rather than treating them as two aspects of the same phenomenon.

This treatment of extinction goes back to the inhibitory language established in the context of Pavlov's conditioning experiments (cf. chapter 2; see also Skinner, 1938, pp. 96-102). That language, developed to deal with combinations of stimulus-presentation and stimulus-control operations, was carried over to consequential operations. Once that had happened, the language was retained because it seemed consistent with effects during extinction other than the defining decrease in responding.

Consider the phenomenon called *spontaneous recovery* (see also chapter 9). In a typical extinction session, responding begins at a relatively high rate that decreases as the session continues. But if the organism is then removed from the situation for some time, the rate at which it begins responding in the next extinction session is usually higher than the rate at the end of the last session. Some hypothetical cumulative records illustrating spontaneous recovery are shown in Figure 4-8. Responding at the beginning of each session is said to have recovered spontaneously from inhibition built up by the end of the previous session; presumably this inhibition builds up within sessions, when it actively suppresses responding, and dissipates during the times between extinction sessions.

Phenomena such as spontaneous recovery were taken to mean that responding that had been reduced by extinction was somehow "there all the time but inhibited." But accounts that assumed responding was actively suppressed in extinction were criticized on the grounds that they explained extinction in terms of unobservable events or processes (R. L. Reid, 1958). For example, when it was said that a response had been inhibited in extinction, the inhibited response could be measured but not any events or processes that were doing the inhibiting. The language of inhibition was not particularly useful if the inhibitors could not be specified. Perhaps for this very reason a variety of competing accounts of extinction were formulated in terms of such inferred processes as inhibition, frustration, interference, and fatigue (cf. Kimble, 1961, chapter 10). These accounts agreed that responding did not merely decline passively in extinction but rather was actively suppressed; they differed only in how they characterized the events that generated the inhibition.

Other accounts, however, did not assume suppressive or inhibitory processes in extinction. For example, one account of spontaneous recovery notes that stimulus properties of the experimental setting may differ for the organism from the beginning to later in the session, perhaps because of lasting effects of handling and other pre-experimental conditions. Thus, the effects of extinction late in one session might not transfer to the different and perhaps novel conditions at the beginning of the next session. On this basis, Kendall (1965) reasoned that the usual pattern of response rates in sessions of extinction should be reversible under proper conditions. In an experiment with three pigeons, he first reinforced key pecks in one-hour sessions. Then he arranged repeated one-minute sessions of extinction. Only after responding had reliably decreased to zero in these brief sessions did he introduce the first long extinction session. Within a few minutes, each pigeon began to respond. Until this session, responding had never been permitted to extinguish at times later than one minute into a session; responding occurred at these later times when the opportunity was finally available. In a sense, Kendall had demonstrated spontaneous recovery within rather than at the beginning of a session.

Furthermore, the temporal pattern of responding in spontaneous recovery




Figure 4-8. Hypothetical cumulative records of spontaneous recovery in successive daily extinction sessions for a rat's lever presses previously reinforced with food. Responding in Session 2 begins at a rate higher than that at the end of Session 1; similarly, responding in Session 3 begins at a rate higher than that at the end of Session 2.

(Figure 4-8) also occurs with responding that has never been explicitly reinforced. With rats newly placed in a chamber with a lever, some lever pressing usually occurs even without food reinforcement; the rate of this *baseline* or *operant-level* responding tends to decrease within sessions and to recover between sessions (Schoenfeld, Antonitis, and Bersh, 1950). Thus, an inhibitory theory would have to argue that the similar effects are only coincidental or to invent some process that could deal with why both types of responding change similarly over time.

Contingencies and stimulus deliveries. Nevertheless, accounts of extinction remain controversial. This may be because discontinuing reinforcement has two effects: a contingency between responses and reinforcers ends, and reinforcers are no longer delivered. The term *contingency* is simply a name for the consequences of responding; it is defined as *the effect of a response on the probability of a stimulus*. For example, if a rat can receive food pellets only by pressing a lever and each lever press produces a pellet, the contingency is one in which the lever press raises the probability of pellet deliveries from zero to 1.0 (no pellet if no lever press and pellet if lever press). If the lever press does not produce pellets, the contingency is one in which pellet deliveries are independent of lever presses (the lever press has no effect on the probability of pellet deliveries).

Let us now compare procedures in terms of changes in contingencies and changes in stimuli. Consider first the food-deprived rat whose lever presses are reinforced with food pellets. Once every ten or fifteen seconds the rat presses the lever and eats the pellet delivered. If lever pressing is then extinguished, lever presses no longer produce pellets and the rat no longer eats. As a result, the rat presses the lever less often. This rat is now different in two ways: not only are its lever presses no longer having their earlier consequences but this rat is no longer eating.

Here is an alternative arrangement. The rat remains food-deprived, but when we discontinue the reinforcement of lever presses we begin to deliver pellets automatically every ten or fifteen seconds. In this case, we change contingencies as in the previous example: lever presses that once raised the likelihood of a pellet delivery now have no effect. But even though lever presses no longer have their earlier consequences, this rat can still eat a food pellet every ten or fifteen seconds. This rat too will press the lever less often; on the other hand, it will differ from the first rat in some important ways.

These two procedures distinguish the effects of terminating a reinforcement contingency from those of terminating reinforcer deliveries. The standard extinction procedure terminates both a contingency and stimulus presentations; in the last example, however, the contingency was terminated without also terminating stimulus presentations. In both cases, the rate of the response that had been reinforced decreases. But terminating the delivery of reinforcers, which occurs only in the first procedure, affects a range of behavior broader than the reinforced response, and this effect is independent of whether reinforcer deliveries had depended on responses. For example, if a food-deprived rat has been eating food pellets and the pellets are suddenly taken away, the rat will become more active, will perhaps urinate or defecate, and may bite or otherwise attack various objects

in the chamber. If pellets were produced by lever presses, the rat might bite the lever (Mowrer and Jones, 1943); other organisms in the chamber might be attacked (Azrin, Hutchinson, and Hake, 1966); and once the procedure makes such aggressive responses probable, the opportunity to engage in them may reinforce other responses (e.g., the organism might pull a chain if chain pulls produced something it could sink its teeth into: Azrin, Hutchinson, and McLaughlin, 1965). These effects may be observed in extinction, but are not direct results of terminating the reinforcement contingency. They are side effects usually superimposed upon the decrease in responding during extinction because extinction is necessarily accompanied by the termination of reinforcer deliveries.

We have compared extinction, in which both a reinforcement contingency and reinforcer deliveries are terminated, with a procedure in which only the reinforcement contingency is terminated. Terminating reinforcement deliveries while maintaining the reinforcement contingency does not seem logically possible. If instead response-independent reinforcer deliveries were terminated, we might expect the same kinds of emotional or aggressive responses as are often produced by extinction; this finding would add to the argument that such changes in behavior are side effects of changes in contingencies.

In the analysis of behavior, operations often have more than one effect. A stimulus that is a reinforcer in a given procedure may affect other responses besides the one that it reinforces: it may elicit some responses and serve as a discriminative stimulus for still others. In the history of research on extinction, the various effects of terminating reinforcer deliveries may have contributed to its treatment in terms of inhibition. Yet the various phenomena seen as indicating that extinction was more than simply the temporary effect of reinforcement were probably only side effects. Many of these phenomena could have been observed in situations that did not involve the consequences of responding (e.g., the aggressive responding generated by terminating reinforcer deliveries). To that extent they are peripheral to an account of the effects of contingencies.

Extinction and superstition. In extinction, the effects of changes in contingencies are confounded with the effects of terminating reinforcer deliveries, because the general effects of terminating reinforcer deliveries are superimposed on the more specific effects of terminating contingencies. Why then has extinction for so long been the primary basis for studying the effects of terminating contingencies? Part of the answer is that disconnecting the lever from the equipment that operates the feeder is more convenient than disconnecting the lever and at the same time substituting a clock that will operate the feeder periodically. Yet it is not likely that this change in apparatus by itself discouraged further research. Instead, the difficulties rested with some additional properties of behavior.

As we saw in chapter 3, the presentation of stimuli can have many effects. A stimulus presented independently of behavior may follow as well as be followed by responses. For example, suppose that the response-independent delivery of a food pellet occurs just after the rat has pressed the lever. What could distinguish this succession of lever press and pellet from that in which the lever press actually produces the pellet? Why should this accidental succession of response and rein-

forcer not affect the rat's future responding just as much as if it were the outcome of a reinforcement operation?

In chapter 3, we discussed *superstition* (Skinner, 1948). Food was presented to a food-deprived pigeon at fifteen-second intervals and responses that occurred just before food deliveries were likely to be repeated (and therefore followed closely by still more food deliveries). Through the accidental succession of responses and reinforcers, stereotyped patterns of responding tended to develop. The responses were different for different pigeons (variability) and they usually changed gradually in form over successive food deliveries (topographical drift). A particular response is not likely to be superstitiously maintained over an extended period of time through the accidental correlation of responses and reinforcers, but the possible effects of such accidental correlations must be considered in any experiment involving response-independent reinforcer deliveries.

The maintenance of superstitious responding will depend to some extent on the temporal spacing of reinforcer deliveries. The optimal spacing is one in which a response made more probable by one reinforcer will be likely to occur again just before the next. Eliciting effects of the reinforcer may also contribute, because the reinforcer may directly generate responses between reinforcer deliveries (e.g., as when pigeons' pecks are made more likely by presentations of grain: cf. chapter 3; Staddon and Simmelhag, 1971; see also Killeen, 1975).

Superstitious responding generated by an accidental succession of responses and reinforcers is a recurrent problem in the analysis of behavior, because such accidental sequences can occur whether reinforcers are independent of or consequences of responses. For example, if one response is followed by a different response that is reinforced, the reinforcer may affect both responses even though its delivery depended only on the second response. Even when a response does have consequences, properties unrelated to reinforcement may become stereotyped if they typically accompany reinforced responses (e.g., the bowler's gestures after releasing the ball may persist because of their close relation to earlier responses and to the impact of the ball; cf. Herrnstein, 1966; see also Guthrie, 1935; Guthrie and Horton, 1946). But the greatest difficulty created by the phenomenon of superstition is that it is too easily invoked to explain behavior not accounted for in other ways. We must avoid abusing the concept while remaining alert to the circumstances in which it may legitimately be applied.

To illustrate the point, reconsider what happens when a reinforcement contingency is terminated while reinforcer deliveries continue. First, a rat's lever presses are reinforced with food pellets; then, when presses no longer produce pellets, pellet deliveries continue independently of behavior. Lever pressing will be maintained for some time after this change in contingencies, and therefore many presses are likely to be followed closely by food pellets. Lever pressing will probably decline only slowly, not necessarily because the reinforcement contingency has enduring effects but rather because the accidental succession of responses and reinforcers may counteract the effects of terminating the reinforcement contingency. Eventually lever pressing will be replaced by other responses (cf. chapter 8), but it would be difficult to argue that such a decrease in lever pressing is simply determined. For this reason, the transition from a reinforcement

contingency to response-independent reinforcer deliveries by itself cannot demonstrate the effects of terminating the reinforcement contingency (cf. Rescorla and Skucy, 1969; Boakes, 1973).

When responses have consequences, the termination of these consequences affects both how stimuli are related to behavior and how stimuli are distributed in time. Each effect must be taken into account in an analysis of behavior. The most important property of extinction is its demonstration that reinforcement has temporary rather than permanent effects, but other events also accompany extinction and can be superimposed on the decrease in previously reinforced responding. It is difficult to conceive of reinforcement without extinction; the effects of extinction tell us that organisms are sensitive to the consequences of their behavior not only when these consequences start but also when they stop.

THE ACQUISITION OF BEHAVIOR

Let us return now to learning, by examining how an organism may acquire responses through reinforcement. Record *A* of Figure 4-9 shows a hypothetical responses through reinforcement. Record *A* of Figure 4-9 shows a hypothetical cumulative record of a rat's very first session of reinforced lever pressing. The first few responses are separated by long pauses. After a little more than five minutes, the long pauses disappear and responding then increases for the rest of the session. The acquisition of lever pressing appears gradual. If we wished to repeat these observations, we might extinguish lever presses until responding decreased to previous levels and then conduct another session of reinforcement. Record *B* of Figure 4-9 shows what this performance might look like. At the beginning of the session, because of the prior extinction, no responding occurs. Finally a response does occur (and is reinforced), and responding immediately rises to a rate approximately equal to the rate at the end of the first reinforcement session. In this instance, we observe an abrupt rather than a gradual acquisition of lever pressing. How shall we reconcile these two very different performances?




Figure 4-9. Hypothetical cumulative records of a rat's initial acquisition of lever pressing during the first session in which each response is reinforced with food (*A*) and reacquisition of lever pressing when reinforcement is introduced again after a period of extinction (*B*).

Sidman (1960) has discussed an alternative to the view that this depends on an irreversible change produced by initial learning:

The animal learned not only those responses that succeeded in depressing the lever but also learned to go to the tray, to pick up the small pellet, to bring it to its mouth, etc. And these responses were learned in their correct sequence, because their reinforcement was correlated with the appropriate stimuli both from the environment and from the preceding behavior. The tray approach, for example, could be reinforced only after the sound from the food magazine; reaching for the pellet could be reinforced only after the pellet had dropped into the tray, etc. . . . What did we extinguish when we disconnected the feeding mechanism? . . . There is no magazine sound, no pellet sound, no visual pellet, no tactual pellet, etc. Tray-approach is still possible, but only in the absence of some of its controlling stimuli. The responses involved in picking up and ingesting the pellet can no longer occur in their originally learned context. While our extinction procedure may have returned the lever-pressing response to its preconditioning level, other components of the total learned sequence could not have undergone complete extinction. (Sidman, 1960, pp. 101-103)

From this analysis, Sidman could summarize the reasons for the difference in acquisition in the two reinforcement sessions as follows: "When reinforcement was again introduced. . . , the animal did not have to relearn the whole sequence, because the whole sequence had not been extinguished" (Sidman, 1960, p. 103).

Reinforcement, then, does not produce learning; it produces behavior. When we look to see whether a rat presses a lever when a reinforcement contingency operates and does not press the lever otherwise, we are simply concerned with the extent to which the rat has learned the consequences of its lever pressing. The consequences of responding are critical to our understanding of learning *not because learning follows from them but because they are what is learned*. Contingencies describe how the environment is affected by behavior and are therefore important features of the environment for organisms to learn.

Latent learning. The preceding issues were implicit in a controversy based upon a phenomenon called *latent learning* (Blodgett, 1929; Thistlethwaite, 1951). Consider the experiment illustrated in Figure 4-10 (Tolman and Honzik, 1930; see also Tolman, 1948). Food-deprived rats in each of three groups negotiated a maze. In one group, the rats found food in the goalbox of the maze, and over successive daily trials their entries into blind alleys gradually decreased. In a second group, the rats found no food in the goalbox. For these animals, entries into blind alleys decreased but remained substantially higher than those for the first group. In a third group, food was introduced in the goalbox only after ten sessions. The performance of this group, which had been like the second group (no food), quickly became comparable to that of the first group (food); the rats that previously negotiated the maze without food in the goalbox now began to run with as few entries into blind alleys as did the rats that had always found food there. Until food was introduced, the learning of the third group had been latent; what had been learned was demonstrated by introducing food.

The argument was that rats learned the maze equally well whether or not




Figure 4-10. An experiment in latent learning. Rats were allowed one trial per day in a fourteen-choice-point maze. One group (filled squares) always found food in the goalbox. A second group (unfilled triangles) never did. The third group found no food in the goalbox through Day 10 (unfilled circles), but thereafter found food (filled circles). This group, which had been performing like the second group, quickly became like the first group. The argument was that the rats in this group had been learning the maze pattern all along, and that food in the goalbox was necessary only to get them to demonstrate what they had already learned. (From Tolman and Honzik, 1930)

food was in the goalbox, and thus learning could not be attributed to food as a reinforcer. The argument was then extended to reinforcers in general, and experiments on latent learning were said to demonstrate that learning could occur without reinforcement. But soon came the counter-argument that food in the goalbox is not the only possible reinforcer for a rat's running of a maze. The rat's removal from the maze at the end of a run, its escape from the constrained spaces in blind alleys, or its return to the home cage where it has been fed all might function as effective reinforcers. Experiments were then performed that varied the rat's handling at the end of the run, the width of the maze alleys, and the rat's home-cage feeding. Each time one experiment showed that a particular reinforcer could generate maze learning, some other experiment demonstrated latent learning in such a way that the reinforcer at issue could not have been effective. And so it went.

Even in principle the argument could not be resolved and latent learning gradually faded away as a critical experimental issue. The reason was that a rat's

negotiation of a maze inevitably involves the consequences of responding. At any choice point, one turn is followed by an entry into a blind alley and the other by an opportunity to move further through the maze; at the final choice point, only one turn is followed by entry into the goalbox, whether or not the goalbox contains food. The rat's sniffing, touching, looking, and moving through the maze are consequential responses, even if this behavior is not as easily accessible to us as correct turns or entries into blind alleys. These consequences are what the rat learns, and whether to call them reinforcers is perhaps mostly a matter of preference.

Sensory-motor learning. The consequences so far considered have mainly been major classes of stimuli, typically with biological significance for the organism's survival (e.g., food and water). But many presumably lesser consequences are important in our day-to-day interactions with our environment. When we open a book, a consequence is that we can read the text. When we listen to someone speaking, a consequence is that we hear what is said. When we reach out for a nearby object, a consequence is that we touch it. Each consequence may then set the occasion for new responses. When we finish one page of a book, we may turn the page and read the next; when the speaker has finished, we may ask a question or make a comment; when we touch the object, we may pick it up and do things with it. To the extent that behavior in each of these circumstances continues because of its consequences, these cases may be discussed in the vocabulary of reinforcement (cf. Skinner, 1953). We may say that seeing reinforces looking, that hearing reinforces listening, and that touching reinforces reaching out (the several responses each occur in appropriate circumstances, but stimulus control is not critical to the present discussion; cf. chapter 7).

The interaction of sensory processes with behavior has been a long-standing source of controversy in the psychology of learning. Theorists took sides in debates on whether learning was sensory or motor. The issues were expressed in various ways. Did organisms learn responses or did they learn relations among stimuli? Did organisms learn response-stimulus or stimulus-stimulus associations?

One problem is whether sensory processes should be treated as behavior. Such a treatment is consistent with the view that behavior must be dealt with in terms of relations among stimuli and responses rather than in terms of stimuli alone or responses alone. Although we cannot measure seeing and hearing as unambiguously as we can measure discrete responses like a rat's lever press or a pigeon's key peck, they are nonetheless components of behavior. They depend not only on whether visual or auditory stimuli are present but also on what the organism does. Looking makes seeing more likely and listening makes hearing more likely. Even with respect to such basic events, the organism is not passive in its environment.

Sounds and lights and other relatively simple environmental events have sometimes been described as neutral stimuli relative to potentially strong reinforcers (e.g., food) or strong punishers (e.g., electric shock). But the label *neutral*, though convenient, is a misnomer. Events cannot be truly neutral if they are consequences of behavior, because it is unlikely that they will be completely without effect on behavior. Nevertheless, in the days before the relativity of reinforcers was appreciated, demonstrations of the reinforcing effects of stimuli

such as lights and sounds were received with skepticism. With repeated experiments, however, the phenomenon called *sensory reinforcement* became established (Kish, 1966). For example, rats' lever pressing in darkness increased if lever pressing briefly turned on a light. Although the effect on lever pressing was small and transient, the light had served temporarily as a weak reinforcer. Phenomena like these soon came to be discussed in terms of *exploratory behavior* and *curiosity* (e.g., Berlyne, 1960), and experiments were extended to a variety of sensory consequences. For example, if a monkey is alone in an enclosed chamber, the opportunity to look at other monkeys on the outside can be used to reinforce the operation of a switch (Butler, 1957).

In these experiments, the experimenter assesses the effects of sensory consequences on a response chosen for its ease of measurement. But in any environment the organism's behavior will have inevitable sensory consequences. The organism changes the part of the environment with which it comes in contact simply by moving from one place to another; the things it sees and touches change as it moves.

An experiment by Held and Hein (1963) illustrates these relations between behavior and simple sensory consequences. Pairs of kittens were raised in darkness; their first experiences with visual stimulation were in the apparatus shown in Figure 4-11. Both kittens received the same kind of visual stimulation: each wore a shield that prevented it from seeing its own feet and body, each was prevented from seeing the other by the wide central post, and each saw the same pattern of black and white vertical stripes that uniformly covered the walls of their circular enclosure. The kittens were linked by a sort of miniature carousel, but one kitten moved actively (*A*) whereas the other kitten was moved passively (*P*). The active kitten stood on the floor of the enclosure whereas the passive kitten stood inside a carrier suspended a short distance above the floor. As the active kitten walked around the central post, the passive kitten in its carrier moved a corresponding distance on the other side, and if the active kitten turned around in place to walk in the other direction, the pulley system turned the passive kitten's carrier so that it too turned to face in the new direction. Both kittens were exposed to the same kinds of visual stimulation in the same temporal sequences, but the visual stimuli for the active kitten were consequences of its own behavior whereas those for the passive kitten were not; they depended on the active kitten's movements rather than on its own. The two kittens were then given standard tests of visual-motor coordination such as visual paw-placement (normal kittens extend their paws toward a table edge or other horizontal surface if held in the air a short distance away from it). Even though their exposure to visual stimuli was equivalent, only the active kitten responded appropriately in these tests; the passive kitten later became able to respond appropriately, after it was allowed to walk about freely in a lighted room.

This experiment by Held and Hein has much in common with a classic experiment by Stratton (1897), who for eight days wore prisms that inverted and reversed his visual fields. At first his world looked upside-down and backwards and his movements were not coordinated with his surroundings. For example, in walking he looked down to see where he was going, but because of the inverting prisms he found that he was looking at the ceiling instead of the floor. Similarly, he had

difficulty in pointing at or reaching for objects, because things once seen below eye level were now seen above eye level and things on the right were now seen on the left, and vice versa. As time passed, however, coordination improved and Stratton reported that the world no longer even looked so upside-down.

Stratton's findings illustrate again the role of the consequences of behavior. Looking and moving in the visual field have different consequences with and without inverting prisms, and adjustment to the prisms requires that these new consequences be learned. For example, seeing the floor as one walks is an important consequence, but looking down, which produced that consequence in the past, no longer does so when one begins wearing inverting prisms; with the prisms, seeing the floor is instead a consequence of looking up (in these situations, of course, up and down can be defined either relative to the visual field or relative to the body; cf. Harris, 1965). If seeing the floor is a reinforcer while one is walking and one walks while wearing inverting prisms, then seeing the floor will reinforce the response of looking up instead of the response of looking down.

With Held and Hein's kittens and with Stratton's inverted vision, we proceeded from the organism's locomotion in the environment to responses of smaller magnitude, such as eye movements. Even such small-scale responses as eye movements can have profound consequences. If we notice something out of the corner of an eye, we are more likely to see it clearly if we look toward it than if we look away from it (except in dim light, when we see an object best only if we look a little bit away from it). Suppose then that, relative to a uniform visual field,




Figure 4-11. An apparatus for studying the relation between behavior and visual stimulation. Both the active kitten (*A*) and the passive kitten (*P*) are harnessed to the pulley system with its fulcrum at the central column. Kitten *A* stands on the platform whereas Kitten *P* rides in a carrier. The changes in position produced by the system (see arrows). (From Held and Hein, 1963, Figure 1)

seeing a contour such as the edge of an object is a reinforcing consequence for eye movement. We should expect eye movements to become coordinated with events in the visual field. Available data from the eye movements of newborn infants are consistent with this view. For example, when shown a simple figure such as a triangle in a uniform visual field, infants tend to fixate more and more accurately on the edges and vertices of the triangle as time passes (e.g., Salapatek and Kessen, 1966; cf. Girgus, 1976; Aslin, 1977).

We have much to learn about how arbitrary such relations between responses and consequences can be. For example, suppose an optical system could project visual stimuli in an infant's current area of fixation. If we chose stimuli known to be effective as reinforcers for eye movement, we could then present these stimuli as a consequence of various kinds of eye movement. Given a stimulus in the right portion of the visual field, we might let it disappear if the infant looked to the right, but move it to the region of fixation if the infant looked to the left (and vice versa for a stimulus in the left portion of the visual field). In other words, this arrangement would create a world in which the infant could fixate on an object only by looking away from it. If eye movements changed accordingly in this kind of artificial environment, we might conclude that similar processes operate in our natural environment (cf. Schroeder and Holland, 1968).

Our visual system is built so that, in the natural environment, particular responses will automatically have particular consequences. The relations between responses and these consequences are the basis for our acquisition of behavior. In situations in which an organism learns, we cannot often say that the organism has learned either stimuli or responses, but we usually can say that the organism has learned the consequences of its own behavior. Reinforcement does not explain this learning; we cannot appeal to it as a cause of learning. Instead, reinforcement is part of the description of what is learned even with respect to our simplest interactions with events in the world.

SUMMARY

We began with a brief history of the Law of Effect: puzzle boxes, mazes, runways, and operant chambers. Out of these apparatuses and methods, the principle of reinforcement emerged as a descriptive term appropriate when responding increased because of its consequences. The reinforcement relation is relative. A stimulus providing an opportunity to engage in one response can reinforce a less probable response, and thus reinforcement relations can be reversed by changing response probabilities through such procedures as deprivation. The effects of reinforcement are temporary, as demonstrated by the phenomena of extinction. The reinforcement of responses necessarily includes the presentation of stimuli, and extinction is therefore complicated because it involves both terminating a contingency and terminating stimulus presentations. The contributions of contingencies and stimulus presentations to behavior can be analyzed by comparing reinforcement, extinction, and superstition procedures. Phenomena such as latent learning and sensory-motor learning demonstrate that reinforcement is not an explanation of learning; rather, it is part of the description of what is learned. Organisms learn contingencies, i.e., the consequences of their own behavior.

Consequences of Responding: Aversive Control

PUNISHMENT

Comparing reinforcement and punishment

The relativity of punishment

Side effects of punishment

Eliciting effects of punishers

Discriminative effects of punishers

Summary

REINFORCEMENT BY STIMULUS REMOVAL (NEGATIVE REINFORCEMENT)

Escape

Elicited responding and escape

The ambiguity of the distinction between positive and negative reinforcement

Avoidance

Species-specific defense reactions

Two-process theories of avoidance

Extinction

Positive and negative punishment

SUMMARY AND EXTENSION: THE LANGUAGE OF AVERSIVE CONTROL

The ethics of aversive control

Punishment and *reinforcement* have fairly straightforward histories. *Punishment* stems from the Latin *poena*, pain or penalty, and *reinforcement* from the Latin *fortis*, strong, which is related to such structural items as *fort* and *burg*. *Escape* and *aversive* are more interesting. *Escape*, as *ex-*, out of, plus *cappa*, cape, is derived from the Old North French *escaper*, to take off one's cloak or, by extension, to free oneself from restraint. *Aversive* is derived from the Latin *a*, away, plus *vertere*, to turn. *Vertere* has a Germanic relative, the suffix *-ward* or *-wards*, to turn attention to, in the sense of guarding or regarding; this suffix appears in *reward*, an occasional synonym for *reinforcement*.

So far we have emphasized a relation, *reinforcement*, in which the consequences of responding make responding more likely. There also exists a relation, *punishment*, in which the consequences of responding make responding less likely. Furthermore, a stimulus that reinforces when produced by a response may serve a different function when responding removes it: its removal may punish responding. Inversely, a stimulus that punishes when produced by a response may reinforce responding by its removal. For example, money may be used to reinforce behavior, as when a child is paid for the completion of a chore, but its removal may be used to punish behavior, as when the child's allowance is canceled because of a misdeed. Similarly, a painful burn may punish behavior, as when one learns not to touch a pan just taken from a hot oven, but its removal or prevention may reinforce behavior, as when one learns to treat a burn with appropriate medication or to put on a kitchen glove while handling things around a stove.

Except for reinforcement by presenting a stimulus (chapter 4), these relations are often grouped together as instances of aversive control. In other words, *aversive control includes both punishment and reinforcement by the removal or prevention of aversive stimuli*. This chapter first treats punishment, and later turns to reinforcement by stimulus removal or prevention, in escape and avoidance procedures.

PUNISHMENT

As an operation, *punishment* is arranging a consequence of responding that reduces the likelihood of the response. The stimulus arranged as a consequence is called a

punisher. For example, if electric shock is delivered to a rat whenever the rat presses a lever, the lever press is said to be punished and the shock is said to be a punisher, because this operation reduces lever pressing. In these respects, the vocabulary of punishment parallels the vocabulary of reinforcement (cf. Table 4-1).

Like *reinforcement*, the term *punishment* has been applied to both operations and processes. Thus, in some literatures stating that a response was punished may mean either that the response produced a punisher or that responding decreased because it produced a punisher. As with reinforcement, preferred usage will be to restrict *punishment* to the vocabulary of operations and to describe the process directly in terms of changes in responding. Because of varied usages in the literature, however, the alternative usage, of punishment as process, will occasionally be convenient.

The vocabulary of reinforcement is also paralleled in the vocabulary of punishment by the object of the operation: responses, not organisms, are said to be punished. If a rat's lever pressing produces shock and lever pressing decreases, it is appropriate to say that the rat was shocked and that the lever press was punished; it is *not* appropriate to say that the rat was punished. This grammatical distinction can have substantial effects on the precision with which we observe and describe behavior.

Consider a child misbehaving. A parent calls the child and when the child comes administers a spanking. To say simply that the parent punished the child may be convenient, but this usage makes it too easy to omit the responses that might be affected. The immediate consequence of the child's misbehaving was that the parent called; the spanking occurred after the child obeyed the call. Although the child might misbehave less in the future because of the spanking, the spanking may also affect the future likelihood that the child will approach when the parent calls. (But it would be inappropriate to recommend that the parent intent on spanking should go to the child rather than having the child come to the parent; too many better alternatives do not require spanking at all, such as reinforcing responses incompatible with the misbehavior.)

The point is not merely grammatical. The detailed analysis is more likely to be made when the punished response must be stated explicitly (spanking punished the child's approach to the parent) than when a less precise description is regarded as acceptable (spanking punished the child). In endorsing the grammar of reinforcing responses and punishing responses, we need not prejudge how these operations affect behavior; we may assume that these operations can affect other responses besides those on which the operations were performed (e.g., a spanking may elicit crying, even though the punishment operation is independent of the child's crying). Working with a vocabulary of operations that unambiguously states the consequences of behavior does not prevent us from describing such effects.

Comparing reinforcement and punishment

The effect of punishment is simply the opposite of the effect of reinforcement. The relation between the two processes is illustrated in Figure 5-1, which presents hypothetical reinforcement and punishment data. The top graph shows

how a rat's lever pressing might change when the rat's lever pressing is followed first by reinforcement with food pellets and then by extinction. During *baseline* performance, when no consequence is arranged for lever pressing, responding occurs infrequently. When *reinforcement* begins, responding increases over the first few sessions, after which it remains at a fairly stable level. *Extinction* then reduces responding to a level that gradually approaches that during baseline.

The bottom graph of Figure 5-1 shows how lever pressing might change during and after punishment of lever presses with electric shock. Because punishment reduces responding, some responding must exist to begin with or no decrease could be observed. In this example, substantial responding is already maintained by food reinforcement and reinforcement continues throughout all sessions; the




Figure 5-1. Comparing the effects of reinforcement and punishment on the hypothetical lever pressing of a food-deprived rat. The top graph shows unreinforced lever pressing (baseline), the increase when lever presses produce food (reinforcement), and the return to earlier levels when reinforcement is discontinued (extinction). The bottom graph shows responding maintained by reinforcement (baseline), the decrease when response-produced electric shock is superimposed on this performance (punishment), and the return to earlier high levels when punishment is discontinued (recovery). The bottom graph uses responding maintained throughout by reinforcement as the baseline against which to illustrate the effects of punishment because a decrease in responding is not easily seen against a baseline in which responding is already infrequent.

effects of punishment can then be assessed by superimposing the punishment operation on this baseline. The *baseline* shows the responding maintained before lever pressing is punished. When *punishment* begins, lever pressing decreases during the first few sessions to a maintained low level. In *recovery*, punishment is discontinued and responding gradually returns to the level observed during baseline.

The figure suggests considerable symmetry of reinforcement and punishment: the former increases responding whereas the latter decreases it, but their effects continue while the procedures are maintained and disappear after the procedures end (responding returns to earlier levels). Actual rather than hypothetical effects of punishment are shown in Figure 5-2 (Estes, 1955).

The effects in Figure 5-2 are clear. Nevertheless, the effectiveness of punishment has long been a matter of controversy. Punishment was incorporated in Thorndike's early versions of his Law of Effect. Thorndike argued then that behavior could be stamped out by annoyers as well as stamped in by satisfiers. Statements of Thorndike's law that included the punishment component were called the *strong* Law of Effect. Later, Thorndike withdrew the punishment component and the version that remained, which included only the stamping in of behavior by satisfying states of affairs, was called the *weak* Law of Effect. Thorndike reached this position through experiments on human verbal learning based on the assumption that the learner's responses could be reinforced by saying




Figure 5-2. Cumulative record of the effect of punishment superimposed upon a reinforcement baseline. The response was a rat's lever press. Food reinforcement decreased during punishment and the punisher was electric shock. Responding (From Estes, 1944, Figure 10.)

"right" and punished by saying "wrong." In the procedures that Thorndike used, saying "wrong" had minimal effects, and the conclusion transmitted to subsequent investigators was not simply that saying "wrong" to a human learner was ineffective as a punisher but rather that punishment in general could not be demonstrated.

Thorndike's conclusion had sufficient impact that even the data in Figure 5-2 were interpreted to mean that punishment was an ineffective procedure. Based on the recovery of responding after punishment was discontinued, the argument was that punishment could suppress responding only temporarily and therefore was not to be taken seriously as a technique for changing behavior. Yet on the basis of this criterion, reinforcement also would have to be judged ineffective. For some reason, the criteria imposed for the effectiveness of punishment were different from those for the effectiveness of reinforcement. Even though responding was reduced during punishment, investigators tended not to attend to that reduction and instead looked only at the recovered responding after punishment was discontinued.

It is difficult to assess the role of the intellectual climate of those times. Psychologists were in a position to make pronouncements about the efficacy of reward and punishment that could influence child-rearing and educational practices across the nation, and ethical as well as scientific judgment undoubtedly contributed to the conclusions. What follows will provide good reasons for endorsing the conclusion that punishment is undesirable for eliminating behavior and that alternative techniques should be seriously considered whenever available. But if this conclusion remains correct, it will be only because Thorndike and his successors were right for the wrong reasons.

In more recent times, investigators turned again to studying punishment and demonstrated some conditions that modify its effectiveness in suppressing behavior (e.g., Azrin and Holz, 1966; Church, 1963; R. L. Solomon, 1964). For example, the following are some conclusions based on experiments with electric shock as a punisher of a pigeon's food-reinforced key pecking: the more intense and immediate the punisher, the more effective it will be; a punisher introduced at its maximum intensity will suppress responding more effectively than a punisher introduced at low intensity and gradually increased to maximum intensity; and the effectiveness of the punisher may change over extended periods of punishment, as when a punisher of low intensity gradually becomes ineffective after many presentations (cf. Azrin and Holz, 1966, pp. 426-27).

These experiments not only changed the criteria for the effectiveness of punishment, but also examined whether punishment had been judged adequately according to the old criteria. On occasion a single punisher or a few punishers, if sufficiently intense and if a consequence of a weakly maintained response, could make a response disappear for days or months or perhaps even the rest of the organism's life; thus, it was argued that punishment could have permanent effects after all. Such cases are exceptional, and we need not be concerned that recovery may sometimes be so slow that it is to all intents and purposes unobservable. Consider the analogous argument for reinforcement. If someone bought a lottery ticket that won a ten-thousand-dollar prize, we would not be surprised to learn

that the occasional buying of lottery tickets continued throughout the person's lifetime even though it was never again reinforced. We would not reject the existence of extinction on those grounds; we therefore should not reject recovery because of occasional instances of durable punishment effects.

Some incidental features of punishment probably contributed to its unusual treatment in the history of the psychology of learning. We noted that a reduction in responding can be studied only if some responding already exists. A response that is never emitted cannot be punished. Thus, experiments on punishment often proceed by superimposing it on responding maintained by reinforcement. But the effects of punishment may then depend on what maintains the responses. For example, if lever presses are punished with electric shock, food-reinforced lever pressing will probably decrease less if a rat is severely food-deprived than if it is only mildly food-deprived.

Another difficulty is that punishing stimuli are likely to have additional effects that occur whether or not the stimuli are produced by responses. As in reinforcement, the effect of punishment must be shown to depend on the *relation between responses and punishers* (contingency) and not simply on the *delivery of punishers*. For example, electric shock may decrease the rate at which a pigeon pecks a key even if shocks are delivered independently of key pecks. It is necessary to demonstrate that the shocks have greater effect when produced by pecking than when they occur independently of pecking (cf. Azrin, 1956).

The prejudices against recognizing punishment were so strong that effective procedures were sometimes even given a different name, *passive avoidance*. For example, assume a rat is on a platform above an electrified grid. On stepping down onto the grid, the rat is shocked and will be less likely to step down in the future. Saying that stepping down was punished by electric shock was appropriate, but it was also possible to say that the organism was passively avoiding the punisher, by holding back from stepping down. This usage allowed effective punishment procedures to be discussed in the vocabulary of passive avoidance while other procedures that did not reduce responding were used to defend the claim that punishment was ineffective.

Cases in which responding decreases because of its consequences may be common or rare in the behavior of organisms, but it would be surprising if none at all existed. Punishment is a name for a relation between responding and consequences. The issue is mainly when it is appropriate to apply this name. Early in its evolution, the analysis of punishment emphasized the impermanence of its effects. For this reason, punishment long went unacknowledged as a fundamental behavioral operation. But the contemporary psychology of learning does not question the existence of consequences that reduce responding. Except that the effects differ in sign, punishment parallels reinforcement: reinforcement increases reinforced responding and punishment decreases punished responding. Both operations have temporary effects: when they are discontinued, responding returns to earlier levels. Because punishment can modify human behavior, questions inevitably will arise about the ethics of its application. But such questions are not likely to be resolved without an adequate analysis of its properties.

The relativity of punishment

In experiments on punishment, punishers are usually chosen for their reliable effect on a variety of responses, because such stimuli reveal the effects of punishment most clearly. One such punisher is electric shock, which can be measured accurately and can be presented at levels that are effective and yet do not damage tissue. Such stimuli, however, are only extreme instances of punishers, and even stimuli that ordinarily serve as reinforcers can become punishers under appropriate conditions. For example, drugs that at first have reinforcing consequences may later become aversive (as when a dose of LSD produces a bad trip). Like reinforcers, punishers cannot be defined in absolute terms nor can they be specified in terms of any common physical properties. Rather, punishers must be assessed in terms of the relation between the punished response and the responses produced by the punisher.

The Premack principle of reinforcement (chapter 4) stated that an opportunity to engage in a more probable response will reinforce a less probable response. The analysis has also been extended to punishment (Premack, 1971). Let us return to the apparatus that can control a rat's opportunities to run in a running wheel or to drink from a drinking tube. It has been modified by a motor that can either lock the running wheel in position, thereby preventing the rat from running, or rotate it at a fixed speed, thereby forcing the rat to run. Rotating the wheel can then be made a consequence of some response, such as drinking from the drinking tube or pressing a lever.

In this apparatus, depriving the rat of an opportunity to run while giving it free access to water makes running more probable than drinking, and depriving the rat of water while giving it an opportunity to run makes drinking more probable than running. Rotation of the wheel can then be made a consequence of drinking: each time the rat drinks, the wheel begins to turn and the rat is forced to run. When running is more probable than drinking, this operation increases the probability of drinking, and it is appropriate to say that drinking is reinforced by running. But when running is less probable than drinking, this operation has an opposite effect: now the probability of drinking decreases when running is its consequence, and it is appropriate to say that drinking is punished by running. This example illustrates either reinforcement or punishment of a single response (drinking) by a single consequence (enforced running), depending on the relative probabilities of these responses. The stimuli and responses in typical experiments on reinforcement and punishment had been chosen to make those procedures work (e.g., with food-deprived rats, eating is far more probable than lever pressing), and thereby obscured the potential reversibility of consequences as reinforcers and punishers. Responding can be raised or lowered by changing its consequences, and these effects are determined by the behavioral and not the physical properties of the consequences.

Side effects of punishment

As with reinforcers, punishers can have effects independent of their contingent relation to responses. If an organism is shocked or burned or pinched,

some of its responses may have little to do with whether or not these events were brought on by the organism's own behavior. Difficulties arise in analyzing punishment because such effects must be distinguished from those that depend on the relation between responses and their consequences. Some effects of shock may be primarily physiological, as when successive shocks systematically reduce a rat's skin resistance (Campbell and Teghtsoonian, 1958). Depending on the nature of the shock source, the effectiveness of later shocks may then vary with the rat's resistance. Other effects are primarily behavioral, as when the apparatus fails to prevent responses by which the organism can reduce its contact with the shock source (e.g., fur is an insulator, and rats have sometimes minimized effects of shock as a punisher by pressing the lever while lying on their furry backs; Azrin and Holz, 1966). In either case such side effects must be taken into account.

Eliciting effects of punishers. Figure 5-3 is from an experiment (Camp, Raymond, and Church, 1967) that compared the effects of response-produced shock with those of response-independent shock. Lever pressing was maintained by food reinforcement in three groups of rats. Measured against a no-shock control group, the responding of both shock groups decreased, but response-produced shock suppressed responding more than response-independent shock (see also Church, 1969). It is therefore appropriate to say that the response-produced shock was a punisher. Events affect behavior most when behavior in turn can affect those




Figure 5-3. Effects of response-independent and response-produced shock on rats' lever pressing maintained by food reinforcement. At the start of the experiment, the rate of shock in the response-independent-shock group was matched to that in the response-produced-shock group. Response-produced shock reduced responding more than response-independent shock. (From Camp, Raymond, and Church, 1967, Figure 5.)

events (Rachlin, 1967, p. 87). Just as we must distinguish between the effects of reinforcer deliveries and those of the contingent relation between responses and reinforcers, so also must we distinguish between the effects of punisher deliveries and those of the contingent relation between responses and punishers.

These considerations are illustrated by the punishment of two classes of species-specific behavior in the Mongolian gerbil (Walters and Glazer, 1971). Sand-digging, part of the gerbil's nesting behavior, consisted of scooping sand and kicking it backwards; alert posturing, a defense reaction occasioned by sudden or aversive stimuli, consisted of standing erect on the hindpaws with ears cocked. Both responses occurred frequently in the sand box where the experiment was conducted. Because delivering shock was difficult in the sand box, a tone was established as an aversive stimulus by repeatedly pairing it with shock in a different setting. When this stimulus was contingent upon digging, it was an effective punisher: digging decreased when it produced the tone and recovered after the contingency was discontinued. The decrease in digging was accompanied by an increase in alert posturing. But when the tone was contingent upon alert posturing, alert posturing increased when it produced the tone and did not return to earlier levels for several sessions after the contingency was discontinued. The increase in alert posturing was not accompanied by any appreciable change in digging. The tone presumably functioned as a punisher for digging but not for alert posturing because the eliciting effect of the tone on alert posturing was more powerful than its punishing effect. Thus, in punishment as in reinforcement it is important to acknowledge the separate effects of response-stimulus contingencies and stimulus deliveries.

Let us turn to another example (Kelleher and Morse, 1968). Consider a monkey in a restraining chair. Shock electrodes are placed on the monkey's tail, and the monkey sits before a lever. At the end of successive ten-minute intervals, a lever press will deliver shock to the monkey's tail. A few minutes after the monkey has been placed in the chair, it begins to press the lever. Eventually ten minutes elapse and the next press delivers a shock. The monkey briefly jumps and squeals, and for a while does not press the lever. But it does begin pressing again, responding more and more rapidly until it shocks itself at the end of the next ten-minute interval. This performance is repeated over and over throughout extended daily sessions. On days when shock is discontinued, the monkey's lever pressing virtually ceases; when shock is reinstated, lever pressing returns. If the shock level is raised, lever pressing increases; if the shock level is lowered, it decreases. Shocks depend completely on the monkey's behavior; the monkey would not receive shocks if it did not press the lever. Why is it then that the monkey does not simply stop pressing?

Behavior like this can be generated in many ways and is easily maintained (e.g., Morse and Kelleher, 1977, pp. 193-98). The paradox is that the same shock that maintains responding when produced by the first lever press after the end of successive ten-minute intervals will suppress responding if produced instead by every lever press. How do we decide whether the language of punishment is appropriate? If every lever press produces shock, we can call the shock a punisher because

this consequence reduces responding. But if lever presses produce shock only at the end of successive ten-minute intervals, we should call the shock a reinforcer because this consequence generates responding. (It would not help to yield to the temptation to call the monkey a masochist. Masochism is simply a name we use when a stimulus that we believe should be a punisher serves as a reinforcer. The term does not provide an explanation.)

One possibility is that the appropriate comparison should be between response-produced shock and response-independent shock rather than between response-produced shock and no shock at all. We know that electric shock elicits such manual responses as lever pressing in monkeys (chapter 3). If response-independent shock produces more lever pressing than equivalent response-produced shock, we may conclude that the eliciting effects of shock are so powerful that they partially override the punishing effects, and therefore that lever pressing occurs in spite of and not because of the punishment contingency. The case may have important implications. For example, the parent who tries to stop a child from crying by punishing the crying may have trouble because the punisher elicits the very response that the parent is trying to suppress. We would be doing a disservice, however, to suggest that our comparison of eliciting and punishing effects has resolved the maintenance of the monkey's lever pressing by electric shock. Such studies would be unjustified if these paradoxical performances could be resolved so easily and superficially. Although the eliciting effects of shock contribute to these performances, these effects do not seem to account for them fully (Morse and Kelleher, 1977). We do not yet know the long-term implications of such performances for our classification of the events here called reinforcers and punishers. But as long as human behavior also includes problems of self-injury and self-abuse, our concern with such phenomena is justified.

Discriminative effects of punishers. Another side effect of punishment may come about because punishers can acquire discriminative properties, as when a response is reinforced only when it is also punished. Consider the alternation of two conditions (Holz and Azrin, 1961). In one, a pigeon's key pecks had no consequences; in the other, every peck produced shock and some pecks produced food reinforcers. Pecking was maintained at a low rate when pecks produced no shock, because then they never produced food either; but pecking increased once pecks began to produce shock, because only then did they occasionally produce food. Sample records with food reinforcement completely discontinued are shown for two pigeons in Figure 5-4. With no shock, the rate of pecking was low. When responses began to produce shock, the rate increased. When shock was discontinued, a brief increase in rate (arrows) was followed by a decrease to the earlier low levels.

Again we may ask whether the shocks can be called punishers. In fact, we should conclude from Figure 5-4 that shock was a reinforcer. The main difference here between the shock and other more familiar reinforcers is that the shock acquired its power to reinforce through its relation to food; it would presumably lose the power if the relation were discontinued. Perhaps these kinds of relations are relevant to human behavior. For example, a battered child might provoke a parent to the point of a beating because beatings are often followed by more




Figure 5-4. Discriminative effects of electric shock. Before the session above, two conditions had alternated: the pigeons' pecks produced neither food nor shock, or the pecks always produced shock and occasionally produced food. In the sessions above, no food was presented. For both pigeons, initially low response rates increased when responses began to produce shock and decreased to earlier levels when shock was discontinued. (From Holz and Azrin, 1961, Figure 3.)

attention from a then remorseful parent than ever follows more peaceful parent-child interactions. A parent's attention can be a potent reinforcer and may often override the effects of consequences that would otherwise serve as punishers. Thus, a behavioral analysis may be relevant to such problems as child abuse. Behavior is determined in multiple ways, and the point of a behavioral analysis is to break complex situations down into their simpler components.

Summary

Punishment is the inverse of reinforcement. Reinforcement is defined by increases in consequential responding whereas punishment is defined by decreases. The vocabulary of punishment parallels that of reinforcement: a punisher is a stimulus and punishment is an operation or a process. The effects of punishment are ordinarily temporary; responding usually recovers to earlier levels after punishment is discontinued. In studying punishment, baseline levels of responding must be high enough to make the reduction in responding easily visible. Thus, experiments often superimpose punishment on a baseline of reinforced responding. An analysis of punishment in terms of baseline response probabilities shows that the effectiveness of punishers, like that of reinforcers, is determined by the relation between the punished response and the responses occasioned by the punisher. In punishment, a more probable response forces the organism to engage in a less probable response. The reversibility of reinforcement and punishment implied by

this relativity can be demonstrated experimentally. Punishment effects can be complicated by eliciting or discriminative effects of punishers. The task of an experimental analysis is to separate such side effects from the primary effect of the punisher. These effects occur because punishment necessarily includes both stimulus presentations and a contingency between responses and stimuli; the effects of the stimulus presentations must be distinguished from those of the contingency.

REINFORCEMENT BY STIMULUS REMOVAL (NEGATIVE REINFORCEMENT)

Organisms may terminate as well as produce stimuli. For example, a rat does not ordinarily expose itself to shock, and if shock does occur the rat will get away from it if given the opportunity. Just as presenting an aversive stimulus can be the basis for punishment, terminating or preventing it can be the basis for reinforcement. When a response terminates or prevents an aversive stimulus and the response thereby becomes more probable, the stimulus is called a *negative reinforcer* and the operation is called *negative reinforcement*. Positive and negative reinforcement are therefore distinguished by whether a response produces or removes a stimulus.

Later we will encounter some problems in the vocabulary of positive and negative reinforcement. Here, we mainly acknowledge a vocabulary with substantial precedent. It is now generally agreed that *positive* and *negative*, as modifiers of the term *reinforcement*, should refer to the nature of the consequence produced by responding (whether the response adds something to the environment or takes something away), and that *negative reinforcer* should refer to the stimulus itself and not to its removal (if removal of shock reinforces a rat's lever press, shock is the negative reinforcer and not the shock-free period that follows the response). This vocabulary was established gradually, however (for some time, negative reinforcement was defined as above in some textbooks but as equivalent to punishment in others). As indicated by the etymologies introducing each chapter, language evolves. Thus, the evolving language of reinforcement may eventually make the distinction between positive and negative reinforcement of marginal value (cf. Michael, 1975). Nevertheless, when we do invoke the vocabulary of positive and negative reinforcement and punishment, we will adhere to the dominant contemporary usage, summarized as follows:

- A. *Reinforcement* makes the reinforced response *increase*.
- B. *Punishment* makes the punished response *decrease*.
- C. The modifier *positive* means that the consequence of responding is the *addition* of a stimulus to the organism's environment.
- D. The modifier *negative* means that the consequence of responding is the *subtraction* of a stimulus from the organism's environment.
- E. *Reinforcers* and *punishers* are stimuli and not the absence of stimuli (given that it is possible to make an unambiguous distinction).

The last item is parenthetically qualified because the distinction is sometimes

difficult. For example, is it more appropriate to think of the effects of a traffic ticket in terms of presentation of the paper on which the ticket is written or removal of money from one's funds on payment of the fine? Similarly, is it more appropriate to think of the effects of performing well on an examination in terms of presentation of a high grade, removal of social pressures exerted by parents or teachers or peers, or prevention of lower grades that might have been but were not given?

Escape

Escape procedures are the simplest examples of negative-reinforcement operations: an organism's response *terminates* an aversive stimulus. Escape procedures are contrasted with avoidance procedures, in which the response prevents or delays an aversive stimulus. This vocabulary is consistent with everyday usage: we *escape* from aversive circumstances that already exist, but we *avoid* potential aversive circumstances that have not yet occurred.

Consider a rat in a compartment with an electrified grid floor. The conditions for escape may be arranged simply by constructing the compartment so that the rat can leave it. Movement from one place to another as the escape response is illustrated in Figure 5-5, from a shock-escape experiment with rats in an electrified runway (Fowler and Trapold, 1962). Running speed was fastest when shock was terminated as soon as the rat reached the end of the runway. The longer the delay between reaching the end of the runway and terminating shock, the slower the running. This is one of many examples of quantitative effects of reinforcement. For example, with both positive and negative reinforcement, immediate reinforcement is more effective than delayed reinforcement (*delay parameter*¹) and large reinforcers are more effective than small reinforcers (*magnitude* or *intensity parameter*).

Movement from one place to another is an escape response that is usually easy to establish, but more discrete responses such as a rat's lever presses are often preferred. In the presence of electric shock a rat's press may turn off the shock, or in the presence of a bright light the press may turn off the light (e.g., F. S. Keller, 1941). If in either case responding increases because of its consequences, we may say the response has been reinforced. Let us therefore compare positive and negative reinforcement: in the absence of food, responding that produces food increases; in the presence of shock, responding that removes shock increases. The parallel is straightforward. Yet in research on negative reinforcement, escape procedures have not received as much attention as more complex procedures with aversive stimuli. Research on negative reinforcement is dominated by avoidance, in which aversive stimuli are prevented or delayed by responses that occur in their absence.

¹ A parameter is a variable in a quantitative relation that determines the effect of other variables (e.g., the effects of reinforcement may depend jointly on delay and magnitude of reinforcement, and we would therefore say that both delay and magnitude are parameters of reinforcement); for most purposes in the analysis of behavior, *parameter* may be read as synonymous with *variable*.




Figure 5-5. Relative running speed as a function of the delay of shock termination in runway escape responding of rats. The longer the delay between reaching the end of the runway and the termination of shock, the slower the rats ran. Each data point is based on the last four of 28 escape trials. (From Fowler and Trapold, 1962, Figure 1.)

Elicited responding and escape. The reason for the relative neglect of research on escape is simple: responses easily increased by positive reinforcement, such as rats' lever presses or pigeons' key pecks, are sometimes difficult to increase by negative reinforcement in escape procedures (e.g., Hoffman and Fleshler, 1959). This is at least in part because of the different temporal relations between the reinforced response and other responses in positive and negative reinforcement. The alternative conditions are diagrammed in Figure 5-6.

In positive reinforcement, the reinforcer is absent when the reinforced response must be emitted. After the response, the reinforcer is presented and other responses occasioned by the reinforcer may follow. For example, if a rat's lever press is the reinforced response and food is the reinforcer, food is absent while the rat is pressing; eating does not occur until the press has been emitted and food has been presented. Lever pressing and eating occur at different times, and do not compete directly with each other.

In negative reinforcement, however, the negative reinforcer is present when the reinforced response must be emitted. After the response, the stimulus is removed. For example, if a rat's lever press is the reinforced response and shock is the negative reinforcer, shock is present when the lever press must be emitted. Some of the responses it produces, such as elicited jumping or squealing, or perhaps behavior that reduces contact with the shock source, may be incompatible with lever pressing. As long as shock is present, these responses remain highly probable and may reduce the probability of lever pressing. Once a lever press is emitted,

because it involves adding energy to the environment in the form of heat when a lever press occurs. On the other hand, relative to a neutral temperature, cold may function as a stimulus through its effects on the temperature receptors in the rat's skin. Although cold is the absence of heat, it is certainly a significant event in the organism's environment and may be regarded as aversive. By this interpretation the procedure should be called negative reinforcement, because turning on the heat lamp when a response occurs terminates the effects of cold.

In escape from cold, it is difficult to specify whether reinforcement involves presenting or removing a stimulus, and it is easy to find ambiguities in other familiar instances of reinforcement. For example, we could argue that water as a reinforcer terminates aversive stimuli generated by a dry mouth, or that food as a reinforcer terminates aversive stimuli generated by the depletion of nutrients in the blood stream (cf. Hull, 1943). Why then did the distinction between positive and negative reinforcement ever take on behavioral significance? Let us reconsider that rat in the cold. Before the reinforced lever press occurs, the rat may huddle in a corner and shiver. These responses may reduce the likelihood that the rat will press the lever. Once the lever press is emitted, the heat lamp turns on and these competing responses become less likely, but the rat cannot escape from cold by pressing if it is no longer cold. In the probabilities of the reinforced response and other competing responses before and after reinforcement, this example seems more like escape from shock than like production of food or water (cf. Figure 5-6). It may therefore be called negative reinforcement more appropriately than positive reinforcement. But this account does not eliminate the ambiguity. Reinforcement changes the organism's environment, and inevitably leads to differences in responding before and after the change. At best, we can regard such changes as producing a continuum of effects, ranging from those in which other responses are highly likely to precede and compete with the reinforced response to those in which they do not do so or perhaps even contribute to the probability of the reinforced response.

Avoidance

In avoidance, the aversive stimulus is not present when the reinforced response occurs. For example, consider the following (Sidman, 1953). A rat is placed in a chamber with a lever and a floor grid through which brief shocks may be delivered. Two clocks control the delivery of shocks to the rat. Which clock runs depends on whether the last event in the chamber was a shock or a lever press. One clock times the *shock-shock* or *SS interval*, the time between successive shocks if the rat does not press the lever. Each time a shock is delivered, this clock resets to zero and starts timing a new SS interval. Whenever the rat does press the lever, control shifts to a second clock. This clock times the *response-shock* or *RS interval*, the time by which each lever press postpones the next possible shock. Each lever press while this clock is running resets the clock to zero and starts a new RS interval. In these circumstances, the rat can postpone the shock indefinitely by always lever pressing before the current RS interval ends. An analogy to the RS-interval clock is a parking meter in which each coin resets the meter to some maximum time; putting in a coin before the meter runs out, like the rat's pressing,

can postpone the violation signal indefinitely. If a shock is delivered at the end of an RS-interval, because no response occurred, control switches back to the SS-interval clock. With this procedure, called *Sidman avoidance* or *continuous avoidance*, avoidance responding can be studied independently of escape responding; shock can be prevented by avoidance responses, but once delivered it is so brief that there is little if any opportunity for its termination.

Data are illustrated for one rat's lever pressing in Figure 5-7, which shows rate of responding as a function of RS interval with SS interval as a parameter. Across functions, the RS interval that produced the maximum rate of lever pressing depended on the SS interval. First consider an avoidance schedule with an RS interval of five seconds and an SS interval of two seconds. Any responding at all will reduce shock, and a rate of one response in every four seconds will avoid shock completely. When the RS interval is shorter than the SS interval, however, some patterns of lever pressing can increase rather than decrease shock. For example, consider now an avoidance schedule with an RS interval of two seconds and an SS interval of five seconds. A rat that never lever presses will receive a shock every five seconds. But if the rat presses every three seconds, a shock will be delivered two seconds after each lever press and the rat will therefore receive twenty shocks per minute. Strictly, if this increase in shock rate reduced responding it would be appropriate to say that responding was punished (cf. Sidman, 1953, 1954). Only if the rat's responding is fast enough that two seconds




Figure 5-7. Rate of Rat 46's lever pressing as a function of RS interval, with SS interval as a parameter. The SS interval is the time between shocks if no responding occurs. The RS interval is the time by which each lever press delays the next possible shock. (From Sidman, 1953, Figure 1A.)

never elapse without a lever press will the rat completely avoid shock, but the increase in shock produced by lower rate responding may prevent it from ever attaining such a performance.

Species-specific defense reactions. The advantage of avoidance procedures over escape procedures is that opportunities for the reinforced response occur in the absence of the aversive stimulus. Thus, other responses elicited or occasioned by the aversive stimulus need not continuously compete with avoidance responding. Nevertheless, just as the ease with which different responses can be established as escape responses varies with different organisms, so also does the ease with which different responses can be established as avoidance responses. For example, pigeons avoid shock more readily if they can do so by moving from one side of a chamber to another than if they can do so by pecking (e.g., Macphail, 1968). The demonstration of avoidance responding makes it more difficult to make the case that such differences depend on competition between the reinforced response and other responses occasioned by the aversive stimulus. It has been argued instead that the differences arise because organisms are variously equipped with defense responses that are species-specific and that success with avoidance procedures will depend upon whether the experimenter chooses a response that the organism is already prepared to emit in aversive situations (Bolles, 1970; Seligman, 1970).

Bolles summarizes the argument as follows:

What keeps animals alive in the wild is that they have very effective *innate* defensive reactions which occur when they encounter any kind of new or sudden stimulus. These defensive reactions vary somewhat from species to species, but they generally take one of three forms: Animals generally run or fly away, freeze, or adopt some type of threat, that is, pseudo-aggressive behavior. These defensive reactions are elicited by the appearance of the predator and by the sudden appearance of innocuous objects. These responses are always near threshold so that the animal will take flight, freeze, or threaten whenever any novel stimulus event occurs. It is not necessary that the stimulus event be paired with shock, or pain, or some other unconditioned stimulus. The mouse does not scamper away from the owl because it has learned to escape the painful claws of the enemy; it scampers away from anything happening in its environment, and it does so merely because it is a mouse. The gazelle does not flee from an approaching lion because it has been bitten by lions; it runs away from any large object that approaches it, and it does so because this is one of its species-specific defense reactions. Neither the mouse nor the gazelle can afford to *learn* to avoid; survival is too urgent, the opportunity to learn is too limited, and the parameters of the situation make the necessary learning impossible. (Bolles, 1970, p. 33)

Two-process theories of avoidance. Avoidance, as the quotation suggests, has often been a focus of controversy. Bolles' point concerns the extent to which avoidance behavior is learned. Another issue is specifying what reinforces avoidance responding. When a successful avoidance response occurs, the important consequence is that nothing happens to the organism. How can the absence of an event affect behavior? According to some views, avoidance responding is maintained because the organism is escaping from some properties of the situation that accompanied past aversive stimuli. This view evolved from earlier procedures in

which a warning stimulus preceded shock and the organism could prevent the shock by emitting the avoidance response in the presence of the warning stimulus. Avoidance was most easily established when the avoidance response not only prevented the shock but also terminated the warning stimulus. It was assumed that the warning stimulus acquired aversive properties of its own through its consistent relation to the aversive stimulus, and therefore that this immediate consequence, escape from the warning stimulus, was the effective reinforcer. In fact, one purpose of Sidman's avoidance schedule was to demonstrate avoidance responding without a warning stimulus.

This view, that the warning stimulus acquired aversive properties from its consistent relation to the aversive stimulus and that the termination of the warning stimulus therefore reinforced the avoidance response, was called a two-process theory of avoidance (Mowrer and Lamoreaux, 1946; Kamin, 1956, 1957; see also Hoffman and Fleshler, 1962). The main issue was the nature of the reinforcer. Some theorists were willing to accept the fact of avoidance responding without appealing to some reinforcing event occurring at the moment of the avoidance response. Others felt that it was necessary to specify such an event. When, as in Sidman avoidance, the event could not be located in some warning stimulus in the environment that terminated with an avoidance response, these theorists instead located the event inside the organism. It was argued, for example, that the organism's state just after an avoidance response, when shock was not imminent, was reinforcing relative to its state long after a response, when a shock might occur at any moment. The latter state acquired aversive properties through its close relation to shock and like a warning stimulus was terminated by a response. Because this change of state presumably occurred immediately with each avoidance response, it was assumed that an account based on other events at some time distant from the response was unnecessary.

Debates on the status of such theories of avoidance responding have a complex history (e.g., Anger, 1963; Herrnstein and Hineline, 1966; Herrnstein, 1969; Hineline, 1977), with some even changing sides (e.g., Schoenfeld, 1950, 1969). Sometimes the issues were reduced to experimental questions: For example, if avoidance responding is to be maintained, must the organism be able to reduce the total number of shocks in a session, or is it sufficient for the organism to be able to delay individual shocks even though the same number of shocks is eventually delivered? It turns out that either condition will maintain avoidance responding (Hineline, 1970; E. T. Gardner and Lewis, 1976).

Another issue was whether avoidance responding could be maintained when responding reduced the likelihood of shocks but did not reliably postpone each shock as in Sidman avoidance. Herrnstein and Hineline (1966) arranged an avoidance schedule in which shock was delivered to a rat with a specified probability at the end of every two seconds. Shock occurred with one probability if no lever press occurred and with a different (usually lower) probability if a lever press did occur. For example, in one condition lever presses reduced the probability of shock from 0.3 to 0.1. Avoidance lever pressing was generated by this procedure, but response rates were sufficiently low to allow varied interpretations (e.g., that

responding occurred because each shock simply occasioned or elicited one or two lever presses). More important, even though no consistent temporal relation existed between individual responses and individual shocks, the likelihood of shock increased more rapidly on the average as time elapsed without lever pressing than as time elapsed after a lever press.. Thus, it was possible to argue that the reduction in average aversiveness produced by a response was a sufficient immediate consequence to reinforce avoidance responding. And so it went.

These issues grew out of the difference between *molecular* and *molar* orientations toward the analysis of behavior. The first approach deals with behavior in terms of moment-to-moment sequences of events in a given setting; the second deals with properties that can be measured only over extended time periods. For example, a molecular approach to avoidance examines the individual time intervals that separate individual responses and individual shocks whereas a molar view examines the more general relation between rate of responding and rate of shock over a large sample of responses and shocks (note that rate can be determined only by sampling events over an extended period of time).

These issues have not been and perhaps cannot be resolved. Either molecular or molar properties of behavior may be important in different settings. Just as organisms may come into situations with species-specific responses likely to be engendered by certain events, it is also reasonable to assume that evolution has equipped them with the capacity to respond differentially to a variety of properties of the situations in which they find themselves. Thus, there seems no a priori justification for assuming that an organism whose responding is determined by the molecular properties of one situation (e.g., the consistent temporal relations between responding and shocks established by the RS and SS intervals of Sidman avoidance) would be incapable of responding according to molar properties of another situation (e.g., the consistent overall relations between rate of responding and rate of shock established by a probabilistic avoidance schedule). If this is so, it is not a matter of choosing one or the other approach but rather of deciding which approach is more appropriate to the analysis of any given situation.

Extinction

As with positive reinforcement and punishment, the effects of negative reinforcement are temporary. And as with these other operations, the effects of terminating contingencies between responses and aversive stimuli must be distinguished from the effects of simply terminating the aversive stimuli. In avoidance, turning off the shock source is sometimes considered an extinction operation. If avoidance responding is maintained at such a rate that shocks are rare, the absence of shocks will make little difference and extinction will be lengthy. In fact, one widely acknowledged property of avoidance responding is its persistence even after aversive stimuli are discontinued. (If instead the procedure is shock escape, turning off the shock source will eliminate responding simply because there is no occasion for escape responding in the absence of the aversive stimulus.)

Consider the alternative. In an experiment with food reinforcement, we can arrange extinction either by turning off the feeder or by breaking the connection

between the response and the feeder. Both methods have the same effect: food is no longer delivered. That is not so with shock escape or shock avoidance. If the shock source continues to operate but responses no longer remove or prevent shock, shocks continue to be delivered. In Sidman avoidance, for example, all shocks would be controlled by the clock for the SS interval; responses would no longer operate the clock for the RS interval. This extinction procedure would discontinue the contingency between responses and shock, but if responding had maintained a low shock frequency it would also substantially increase shock deliveries. The effects of changing the rate of shock would have to be separated from those of changing the contingency.

The first method, discontinuing the aversive stimulus, has been the more common extinction procedure in avoidance. Because the absence of the aversive stimulus cannot be a reinforcing consequence unless the aversive stimulus is sometimes presented, however, the second procedure, presenting the aversive stimulus while discontinuing the consequences of responding, is sometimes regarded as more appropriate. In any case, extinction after negative reinforcement shows that the effects of negative reinforcement are temporary, but the time course of extinction depends on the nature of the extinction operation and the rate at which aversive stimuli occur before and after the onset of extinction (cf. Hinde, 1977, pp. 377-81).

Positive and negative punishment

The distinction between positive and negative reinforcement can easily be extended to positive and negative punishment (though here too ambiguous cases are possible). Responses may be punished by some events, such as shock or the rotation of a wheel that forces a rat to run. Responses also may be punished by terminating events. For example, the removal of food as a consequence of a food-deprived rat's lever pressing is likely to reduce lever pressing and therefore is an example of negative punishment. The problem is that punishment might be difficult to demonstrate in this instance. If the rat is food-deprived and food is available, the rat will probably eat rather than press the lever and thus we will have few opportunities to punish lever pressing by removing food (consider Figure 5-6 with food and shock reversed). For this reason, studies of negative punishment have usually removed a stimulus in the presence of which responses are reinforced rather than removing the positive reinforcer itself. (This parallels the emphasis on avoidance rather than escape in studies of negative reinforcement.)

For example, assume that two levers are available to a monkey and that presses on one lever produce food whenever a light is on. We can expect some pressing on the other lever, but this pressing can be punished by making each press produce two minutes during which the light is turned off and presses on the first lever no longer produce food. Such periods are called *time-out*, and the operation is sometimes called *punishment by time-out from positive reinforcement* (e.g., Ferster, 1958; Kaufman and Baron, 1968). Negative punishment has also been called *omission training*, but this name is typically restricted to cases in which the operation is signalled (e.g., Sheffield, 1965).

SUMMARY AND EXTENSION: THE LANGUAGE OF AVERSIVE CONTROL

The presentation or removal of stimuli can reinforce or punish behavior. The empirical findings with reinforcement and punishment suggest that reinforcement is most effective if the reinforced response is compatible with or at least independent of the responding occasioned by the reinforcer, but that punishment is most effective if the punished response is incompatible with or is interfered with by the responding occasioned by the punisher. Thus, it may be easy to reinforce jumping or running with shock removal (escape), but it may be difficult to punish jumping or running with shock presentation.

The stimuli that can reinforce by their presentation can also punish by their removal, and vice versa. We have therefore spoken of *punishers*, of *negative reinforcers*, and of *aversive stimuli*. Each was introduced in a different context, but this was fitting because context determines the behavioral function of any stimulus. Thus, aversive stimuli were introduced in connection with elicitation phenomena, punishers entered with the discussion of consequences that reduced responding, and negative reinforcers joined these other terms as one class of consequences that raised response probability. It would be convenient if we could assume that each term identified different aspects of a single category of environmental events. We might then speak of shock interchangeably as an aversive stimulus, a punisher, or a negative reinforcer, depending on the current situation.

For many stimuli much of the time, this assumption is probably correct. If we know a stimulus is effective as a punisher, we will reasonably expect it also to be effective as a negative reinforcer; this consistency is part of our justification for calling it aversive. Consistencies are to be expected because the categories have their origins in relations among the probabilities of different responses in the behavioral hierarchy. But we must beware of taking this assumption too much for granted. The very fact that we may easily reinforce jumping with shock removal whereas we may not so effectively punish it with shock presentation shows that the symmetry of reinforcement and punishment is limited.

Failures of symmetry between reinforcement and punishment perhaps have encouraged attempts to reduce either one to a special case of the other. Some instances of punishment have been described in the language of *passive avoidance*: by not responding, the organism was passively avoiding the stimulus arranged as a punisher for responding. But then we might as well say that no responding (e.g., not stepping down from a platform onto an electrified grid) is a response that can be reinforced (cf. Dinsmoor, 1954). If the language works in this case, why not in any case of punishment?

Implicit in these accounts is the question of what counts as behavior. If responding is punished, we may say that not-responding is reinforced, and if responding is reinforced, we may say that not-responding is punished. The difference between reinforcement and punishment appears to vanish with such an extension. Yet we know that we usually can tell the difference between instances of reinforcement and instances of punishment. In fact, the differences are often of

serious concern. For example, our society is not indifferent to whether a parent reinforces cooperative behavior with praise or punishes the failure to cooperate with beatings.

We can speak more easily in terms of discrete responses than in terms of the absence of responses, and therefore we prefer direct descriptions in terms of measurable responses such as lever presses or key pecks whenever such descriptions are possible. An organism may exhibit more or less behavior at any given time, and even though we may never observe a complete absence of behavior as long as an organism lives, we need not assume that all failures to act are in themselves actions. Just as we do not have to achieve a temperature of absolute zero to recognize temperature as a dimension of matter, so we do not have to produce a completely nonbehaving live organism to recognize that the quantity of an organism's behavior can change from time to time. The implication is that if we punish a response, we should know what happens to that response before we go looking for an account in terms of other responses, and we should know what else the organism is doing before we begin speaking of not-responding as behavior.

The ethics of aversive control

The behavioral properties of aversive control would lead us often to agree with those who argue against aversive control on ethical grounds. For example, it is important to point out that a parent who arranges aversive contingencies for a child's behavior may acquire aversive properties; to the extent that the child then learns to escape from or avoid the parent's company, contingencies other than those available to the parent are likely to begin to influence the child's behavior. But there are possible exceptions: if punishment seemed the only technique available to reduce the dangerous self-mutilating behavior of an autistic child, punishment might be a lesser evil than the permanent damage the child might self-inflict. To the extent that ethical precepts are concerned with acceptable and unacceptable outcomes of our actions, examining the consequences of responding cannot help but be important.

Skinner outlines our present status in the following quotation; his points are mainly directed to the differences between reinforcement and punishment, but are easily extended to other areas of aversive control:

Civilized man has made some progress in turning from punishment to alternative forms of control. Avenging gods and hell-fire have given way to an emphasis upon heaven and the positive consequences of the good life. In agriculture and industry, fair wages are recognized as an improvement over slavery. The birch rod has made way for the reinforcements naturally accorded the educated man. Even in politics and government the power to punish has been supplemented by a more positive support of the behavior which conforms to the interests of the governing agency. But we are still a long way from exploiting the alternatives, and we are not likely to make any real advance so long as our information about punishment and the alternatives to punishment remains at the level of casual observation. As a consistent picture of the extremely complex consequences of punishment emerges from analytical research, we may gain the confidence and skill needed to design alternative procedures. (Skinner, 1953, pp. 192-93)

Operants: The Structure of Behavior

**SHAPING: THE DIFFERENTIAL REINFORCEMENT OF
SUCCESSIVE APPROXIMATIONS**

OPERANTS: DIFFERENTIATION AND INDUCTION

Response classes

A hypothetical example of differential reinforcement

Differentiable properties of behavior

Schedules of differential reinforcement

Complex behavior: Maze learning

Response sequences: Chaining versus temporally extended units

Limits on differentiation

SUMMARY AND EXTENSION: THE STRUCTURE OF BEHAVIOR

SIX

An operant is a class of responses. *Operant* can be traced to the Latin *opus*, work, which is also a source of *operation* and *copy*. *Class* can be traced to the Latin *classis*, a division of Roman citizens eligible for military draft and perhaps thereby a summons or call. *Class* is related to *claim* and *calendar* through the Indo-European root *kel-*, to shout. It is tempting to assume a relation to the similar root, *gal-*, to call, which gives us *call* and *recall*, but the evidence does not support that assumption.

We have seen how to change behavior by presenting stimuli (elicitation) and by arranging consequences for responding (reinforcement and punishment). One way to discuss the effects of these operations is in terms of the relative positions of responses in the behavioral hierarchy. The ordering of responses in this hierarchy is a property of behavior with which the organism comes to us: the organism's behavior consists of a repertory of responses, each with a different probability. But if we restrict our attention only to these responses, we miss one of the most interesting features of those changes in behavior called learning, i.e., the circumstances in which an organism comes to respond in new ways. In learning, an organism may emit responses it was unable to emit before learning began. We must therefore examine how new responses may be added to an organism's behavioral repertory. In this chapter, we begin by considering *shaping*, a procedure for generating new responses. This will lead us to consider how classes of responses are defined as units of behavior, in the concept of the *operant*. This concept provides the basis for discussing the structure of behavior.

The rat's lever press and the pigeon's key peck have often served as our examples of responses. But if we simply place a rat in front of a lever or a pigeon in front of a key, we may or may not observe a lever press or a key peck. With some

organisms, we will be fortunate and these responses will occur after not too long a wait; with others we may have to wait one or several hours; and an occasional organism may remain so long without responding that our patience is exceeded. Reinforcement cannot have any effect if the response to be reinforced is never emitted. An alternative is available, however. Instead of waiting for the response, the experimenter may generate a response by successively reinforcing other responses that more and more closely approximate it.

SHAPING: THE DIFFERENTIAL REINFORCEMENT OF SUCCESSIVE APPROXIMATIONS

Consider the pigeon's key peck. Once the pigeon begins to eat whenever the feeder is operated, the experimenter operates the feeder only when the pigeon turns toward the key. After reinforcing two or three movements toward the key, the experimenter may then reinforce not just any movement toward the key, but only movements that include forward motion of the pigeon's beak. By this time, the pigeon spends most of its time in front of the key, and the experimenter can shift attention from the pigeon's turning toward the key to the pigeon's forward beak movements. The forward beak movements more closely approximate key pecking than turns toward the key, and once their reinforcement has guaranteed that they will continue to occur, it is no longer necessary to reinforce turning toward the key. By this time, the pigeon's beak movements are full-fledged pecks, and soon one of them strikes the key. At this point, the experimenter can withdraw, because the apparatus can be arranged so that subsequent pecks operate the feeder automatically.

An experimenter experienced with pigeons can usually shape a key peck within ten or fifteen reinforcements. Some aspects of skill in shaping can be explicitly stated. For example, reinforcing a movement is more likely to shape responding efficiently than reinforcing a posture. Other aspects cannot be formulated so readily. For example, shaping must ordinarily compromise between extremes of frequent and infrequent delivery of reinforcers. The frequent delivery of reinforcers will lead to quicker satiation and may overly strengthen some responses that later will not be part of the response that is the end-product of shaping. On the other hand, the infrequent delivery of reinforcers may reduce responding in general, and once the organism becomes inactive the progress in shaping up to that point may be lost. Thus, the experimenter must work within the limits imposed by these extremes; but explicit rules for judging just where these limits will be for any particular organism do not exist (but see Platt, 1973). Furthermore, some features of shaping are fairly specific to the particular response and organism under study, whereas others are generally pertinent to shaping a variety of responses in a variety of organisms. For example, an experimenter who has worked often with pigeons knows that reinforcing a small beak movement aimed directly at the key will more effectively establish key pecking than reinforcing

ing a large sidewise beak movement that finishes in front of the key. On the other hand, whatever the response and the organism, the experimenter should not miss an opportunity to reinforce a response if it more closely approximates the response to be shaped than any response that had been reinforced before. We are not so much concerned with the details of the pigeon's key peck as with the general properties of shaping, which may be applicable to the mastery of such varied skills as gymnastics, handwriting, playing a musical instrument, and pronouncing the sounds of a foreign language.

As a procedure, shaping is based upon *differential reinforcement*: At successive stages, some responses are reinforced and others are not. In addition, as responding changes, the criteria for differential reinforcement change, in *successive approximations* to the response to be shaped. The property of behavior that makes the differential reinforcement of successive approximations to a response effective is that behavior is variable. No two responses are the same, and the reinforcement of one response produces a spectrum of responses each of which differs from the reinforced response along such dimensions as topography (form), force, magnitude, and direction. Of these responses, some will be closer to the response to be shaped than others and may therefore be selected next to be reinforced. Reinforcing these responses will in turn be followed by still other responses some of which may come even closer to the response to be shaped. Thus, reinforcement can be used to change the spectrum of responses until the response to be shaped occurs.

This aspect of shaping, which depends on the relation of a reinforcer to the responding that produced it, is sometimes supplemented by another and more general effect of reinforcers. Some reinforcers, such as food, often increase activity (e.g., the delivery of food to a food-deprived pigeon makes the pigeon active; it is therefore difficult to use food to reinforce the holding of a posture: cf. Blough, 1958). A response more closely approximating the response to be shaped than the one most recently reinforced may occur not only because reinforcement makes such responses more likely but also because some reinforcer deliveries make an organism more active.

There is a paradox to shaping. Reinforcement is said to raise the probability of the reinforced response. But no response is ever repeated exactly. How then can we appeal to reinforcement as the basis for the shaped response when reinforcers are delivered after responses that only approximate it? In fact, if individual responses are never repeated how can we speak of reinforcement at all? We acknowledged this problem in chapter 2 but did not resolve it. In what follows, we will see that we must treat not single responses but rather classes of responses.

OPERANTS: DIFFERENTIATION AND INDUCTION

If we watched a rat's lever presses, we might see the rat press the lever with either paw or with both paws or by sitting on it or perhaps even by biting it. Each is a different response, and even if two presses were made with the same paw they

would not be identical. Nevertheless, we call all of these responses lever presses. On the other hand, the rat might make similar movements at the other end of its chamber, distant from the lever, but we would not call those responses lever presses no matter how closely they resembled the earlier responses that did operate the lever.

It is not sufficient to speak of behavior only in terms of individual responses. Individual responses are instances of behavior and each can occur only once, because responses can have common properties but cannot be identical in all respects. The reinforcement of one response produces later responses resembling the reinforced response more or less closely, but the later responses cannot be exactly the same as the reinforced response. On the other hand, we cannot group all responses together without distinction, because we would be left with nothing to speak of but behavior in general. We must settle for an intermediate level of analysis, at which we speak neither of individual responses nor of behavior in general, but rather of *classes of responses* defined by common properties.

Response classes

In experiments on the rat's lever press, the lever is attached to a switch that operates whenever the rat moves the lever with sufficient force through a sufficient distance. The common property of all lever presses is this consequence of the press: each response that operates the switch qualifies as a lever press. Defining response classes in terms of common environmental effects is the basis both for measuring responses in the class and for imposing experimental operations on these responses. For example, an experimenter could measure lever presses by recording operations of the switch and arrange that all responses with this effect are reinforced with food pellets.

But the response class established by the experimenter has no behavioral significance without taking into account the effects of the experimental operations imposed on this class. The experimenter must ask a fundamental behavioral question: can the consequences of responses in this class modify the likelihood of responses in this class? If so, the class may be called an *operant* class; an operant class is a class of responses affected by the way in which it operates upon or works on the environment. Lever presses and key pecks are convenient examples, but more extensive and complex cases are included among operant classes.

Early in the psychology of learning, when operant behavior was called *instrumental* or *voluntary*, it was assumed that only responses of the skeletal musculature could enter into classes modifiable by their consequences. Other classes of responses called *autonomic*, such as those of glands and smooth muscles, had not been shown to be modifiable by their consequences. Such responses were typically elicited, and Pavlov's conditioning procedures (see chapters 2 and 9) had shown how new stimuli could come to elicit them. At that time, there was negligible evidence that these responses could be modified through reinforcement or punishment.

The demonstration that autonomic responses could sometimes be modified by their consequences changed this view. Consider the salivary response. Salivation occurs spontaneously as well as when elicited by a stimulus such as food in the mouth (e.g., Zener and McCurdy, 1939; cf. chapter 3). Consequences could be arranged for spontaneous or emitted salivation, as measured in drops of saliva brought out from a dog's salivary duct to an external recording system. The consequences of salivation could not be food, however, because the effect of food as a reinforcer would be hard to distinguish from the effect of food as an eliciter of salivation. Water was therefore used to reinforce salivation in water-deprived dogs (N. E. Miller and Carmona, 1967); water does not elicit salivation. Salivation increased when it produced water (reinforcement) and decreased when it prevented the delivery of water (punishment). Responses in this autonomic class, salivation, were modified by their consequences; in other words, these autonomic responses could appropriately be called an operant class.

Another example is illustrated in Figure 6-1. Because autonomic responses can be modified by skeletal responses (as when physical exertion raises heart rate), and because they are assumed typically to vary together (as when excitement or relaxation jointly affect heart rate, breathing, and gastric secretion), a convincing demonstration of operant control over autonomic responses should eliminate skeletal responding and should show the independence of different responses. This was accomplished with groups of rats paralyzed with curare (N.E. Miller and Banuazizi, 1968). Electrical brain stimulation, an event known to reinforce responses as lever pressing in the rat (Olds and Milner, 1954), was used as a reinforcer. This reinforcer was made contingent upon intestinal contraction or intestinal relaxation, and upon heart rate increases or heart rate decreases. Figure 6-1 shows that these two classes of autonomic responses were independently affected by their reinforcers. Some questions have since been raised about whether these and similar studies are reproducible (N. E. Miller, 1978), but the procedure nevertheless remains a methodological *tour de force* in the differentiation of response classes.

An operant is a response class that can be modified by the consequences of the responses in it. This definition of a response class depends on behavioral properties of responding and not on independent physical or physiological properties (such as the somatic-autonomic distinction). These behavioral properties are based on the operation called *differential reinforcement*, the reinforcement only of responses that fall within a specified class. This operation makes subsequent responding conform more and more closely to the defining properties of the class. The essential feature of an operant is the correspondence between a class of responses defined by its consequences and the spectrum of responses generated by these consequences.

A hypothetical example of differential reinforcement

Let us place a food-deprived rat in a chamber with a horizontal slot 30 centimeters or about 12 inches long in one wall (cf. Antonitis, 1951; Herrnstein, 1961b;




Figure 6-1. Differential reinforcement of intestinal motility and heart rate by electrical stimulation of the brain of curarized rats. Intestinal motility changed when intestinal contraction or relaxation was reinforced but not when a change in heart rate was reinforced (top graph). Heart rate changed when fast heart rate or (bottom graph). Curarization, which causes paralysis, eliminated motor responses that might otherwise have been implicated, and the differentiation of the two response classes indicates that the changes were not caused by the general level of autonomic arousal. (From Miller and Banuazizi, 1968, Figures 3 and 4.)

Gollub, 1966; Eckerman and Lanson, 1969). A series of photocells permits us to measure where the rat pokes its nose through the slot. We will label successive 2-centimeter segments along the length of the slot as positions 1 through 15, reading from left to right. Food pellets can be delivered into a cup in the wall opposite the slot. Pellet deliveries are accompanied by a distinctive noise, and the rat quickly comes to the cup and eats whenever a pellet is delivered.

In the absence of reinforcement, the rat spends only a little time near the slot. Occasionally it sniffs at the slot and puts its nose in it as it moves about, but these responses are relatively infrequent and do not have any systematic relation to the positions along the slot. A frequency distribution of the rat's responses as a function of position is shown in *A* of Figure 6-2.

Now let us reinforce the rat's poking of its nose through the slot, but only if the nose-poke occurs at positions 9, 10, 11 or 12. The initial effect of reinforcement, illustrated in *B* of Figure 6-2, is to increase responding at all positions; this effect is not restricted only to those positions correlated with reinforcement. This phenomenon, the spread of the effect of reinforcement to other responses not included in the reinforced class, is called *induction* (an occasional synonym is *response generalization*). In the example, reinforcing responses at positions 9 through 12 affected not only responding at those locations but also responding at other positions across the entire slot.

As we continue differential reinforcement, reinforcing responses at positions 9 through 12 but not at other positions, responding gradually increases at positions correlated with reinforcement whereas it decreases at the other positions. These effects are shown in *C*, *D* and *E* of Figure 6-2. Eventually, most responses are within the boundaries that determine whether a response is to be reinforced, as in *E*, and a point may be reached at which, even though some unreinforced responses still occur, the distribution of responses across positions does not change with continued differential reinforcement.

In the example, the distribution of emitted responses came to conform closely to the boundaries of the class of reinforced responses. This process is called *differentiation*, and responding established in this way is said to be differentiated. Differential reinforcement established a response class defined in terms of response position. Yet if the distribution of responses in *E* represents the maximum differentiation possible with this procedure, what can be said about the responses that continue at positions 6, 7 and 8 or at positions 13, 14 and 15? They are outside the boundaries of the class of responses correlated with reinforcement, and according to a strict interpretation of the defining properties of operants they cannot be counted as members of the operant class. They can at least be spoken of in terms of induction: These responses are so close to the class of reinforced responses along the continuum of position that the effects of reinforcement have spread to them from the reinforced class. This view simply attributes responding within the boundaries of the reinforced class to differentiation and responding outside these boundaries to induction. The same operation generates responding outside as well as inside these boundaries, and this responding can be represented as one continuous distribution.




Figure 6-2. Hypothetical response distributions illustrating the effects of differential reinforcement on response location. A rat pokes its nose into a 30-centimeter horizontal slot in the chamber wall; a photocell system allows this response to register at one of fifteen positions reading from left to right across the slot. The distribution of response positions when no responses are reinforced is shown in *A*. In *B* through *E*, responses at positions 9 through 12 (bounded by dashed vertical lines) are reinforced with food pellets; filled areas represent reinforced responding. In *B*, the effects of reinforcement spread to responses across the entire length of the slot; this spread of the effect of reinforcement is called *induction*. In *C* through *E*, with reinforcement; this narrowing or concentration of the effect of reinforcement is called *differentiation*. Thus, induction may be regarded simply as the inverse of differentiation. With continued differential reinforcement, the distribution of responses becomes stable and corresponds fairly closely, as in *E*, to the class of responses correlated with reinforcement.

The difficulty can be resolved by recognizing that this example involves two different classes. The first class is the basis for the operation of reinforcement and is represented by dashed vertical lines in Figure 6-2; these lines show how stimuli (food pellets) depend upon response position. The lines may be thought of as representing a distribution of stimulus probabilities as a function of position; the probability of a pellet is 1.0 following responses in positions 9 through 12 and is zero elsewhere. This distribution defines a class of responses in terms of the consequences of responses within that class. The second class of responses is given by the actual performance resulting from the reinforcement operation. This class is represented at any time by the current distribution of responses (e.g., early in the course of differential reinforcement, as at *C*, or late, as at *E*). The two classes need not correspond exactly. In fact, the degree of correspondence between the behavior that is reinforced and the behavior generated by this reinforcement is a fundamental dimension of any class of reinforced responses.

Let us now consider another hypothetical example, illustrated in Figure 6-3 (cf. Herrick, 1964). Again we use photocells to record the positions at which the rat pokes its nose through a slot, but this time we use a vertical rather than a horizontal slot. The fifteen positions are numbered consecutively reading from the bottom to the top. Before reinforcement begins, the rat occasionally puts its nose in the slot as it sniffs about the chamber, but these responses occur predominantly at the lower positions, as in *A* of Figure 6-3. Perhaps a response would eventually occur at position 15, at the top of the slot, if we were patient enough. But if we are interested in differentiating responding at the upper positions of the slot, shaping is a preferable option.

Reinforcement is first arranged for responses at position 7 or higher, as in *B*. Responding increases, but most responses continue at the lower positions. Later, the criterion for differential reinforcement is raised to position 9, as in *C*. By this time, responding has decreased at the lower positions, the maximum of the distribution has shifted toward the higher positions, and responding has occurred for the first time at position 13. Responding becomes more concentrated in the region correlated with reinforcement in *D*, when shaping continues by raising the criterion to position 11. Finally, in *E*, the criterion is raised to position 13. Had this criterion been imposed on the initial performance, in *A*, no reinforceable responses might ever have occurred. With the differential reinforcement of successive approximations, however, the response distribution has been shifted to higher levels, with maximum responding at position 12. Nevertheless, this maximum remains below the boundary that separates reinforced from unreinforced responses, and more responses are unreinforced than reinforced in these circumstances. The basis for this outcome is simple: the rat can reach some positions more easily than others. But what implications does such an outcome have for defining response classes in terms of their consequences?

In this instance, the class of responses defined by consequences (responses at positions 13 and higher, which were reinforced) differed from the class of responses generated by the consequences (responding that extended up to position




Figure 6-3. Hypothetical response distributions illustrating the effects on response location of differential reinforcement of successive approximations to a response (shaping). The situation is similar to that in Figure 6-2, except that the slot is vertical instead of horizontal, and the fifteen positions are numbered from the bottom to the top of the slot. As in Figure 6-2, unfilled areas represent unreinforced responding and filled areas represent reinforced responding. Before reinforcement, in *A*, more responding occurs at the lower than at the higher positions along the slot. In *B*, responses at position 7 or higher are reinforced; the lower boundary of the reinforced responses is shown by the dashed vertical line. After this differential reinforcement has some effect, the criterion position is moved up to position 9, in *C*. Later, it is moved to position 11, in *D*, and finally to position 13, in *E*, which shows stable performance after continued differential reinforcement with this criterion. Shaping produces some responses at position 15, the top of the slot, that had not been observed before reinforcement. Nevertheless, the distribution of responding that is maintained at this point has its maximum at position 12, just below the minimum position required for reinforcement. In this instance the distribution of responses corresponds only in part to the class of responses correlated with reinforcement.

15, but with the maximum at position 12). Responding was obviously modified by its consequences, but some of this responding was outside the boundaries of the criterion for reinforcement. Thus, the operant class must be defined in terms of the relation between the environment (the consequences that it arranges for responses) and behavior (the responding produced by these consequences). For convenience, we will occasionally speak of an operant solely in terms of a class defined by consequences or solely in terms of the distribution generated by these consequences, but it is important to remember that the operant is more strictly defined in terms of the correspondence between these two classes (cf. Catania, 1973a).

Differentiable properties of behavior

Reinforcement inevitably includes differentiation. Responses such as lever presses and key pecks must occur at a particular location and be of sufficient force or they will not produce the consequences arranged for them. We could repeat the examples of Figures 6-2 and 6-3 by substituting such dimensions as force of a lever press or a key peck. Responding can vary not only in its location or its force but also in its topography or form, its duration, its direction, and so on. Thus, because differential reinforcement can be based on any dimension of responding, any dimension may provide the defining properties of an operant class.

Figure 6-4 shows a case of response differentiation in a human (Hefferline and Keenan, 1963). The response was an invisibly small twitch of the thumb, measured in microvolts of muscle contraction. The reinforcer for thumb contractions in the 25 to 30 microvolt range was the addition of counts each worth a nickel to a visual display. Relative to baseline responding, in which thumb contractions in the reinforced range were infrequent, responses in the reinforced class interval together with those in the adjoining class intervals increased whereas responses in the 10 to 20 microvolt range decreased. Induction to higher voltage contractions was also evident. In extinction, the form of the distribution began to return toward that of the baseline distribution. For each participant this procedure differentiated thumb contractions in the range correlated with reinforcement, but none was able to report on the thumb twitching or on what had produced the reinforcers. This experiment has been cited to defend the argument that the effects of reinforcement do not depend on an organism's awareness of the reinforced response or the reinforcement contingency.

Temporal properties of responding include latency, duration, and rhythm, and any of these properties may provide the basis for differential reinforcement. For example, Figure 6-5 shows data from a procedure in which a pigeon's peck on a lit key was reinforced with food only if some minimum time had elapsed since the key light turned on; in other words, the procedure differentially reinforced long latencies (Catania, 1970). Each distribution in the figure was obtained at a different criterion latency for reinforcement, represented by the dashed line. The shaded areas represent responding with latencies too short to qualify for reinforcement, and the unshaded areas responding with latencies long enough to qualify for




Figure 6-4. Differentiation of a small-scale response in a human. Invisibly small thumb contractions were recorded in five-microvolt class intervals. The threshold range from 25 to 30 microvolts (filled bars) was ten microvolts. Contractions in the worth nickels to a visual display in front of the participant. The frequency histograms show the performance of one of four humans over consecutive ten-minute segments of a single 80-minute session. (From Hefferline and Keenan, 1963, Figure 3.)




Figure 6-5. Latency distributions of Pigeon P401's key pecks (Catania, 1970). Latencies were timed from the onset of the key light, and a peck was reinforced with food if its latency exceeded some criterion value. The distributions were obtained at each of several latency criteria, ranging from 0.6 to 48 seconds. Filled areas represent unreinforced pecking and unfilled areas represent reinforced pecking.

reinforcement. When the criterion latency was ten seconds or less, the peak of the distribution was located at a value longer than the criterion latency and more pecks were reinforced than unreinforced (distributions at the front left). When the criterion latency became longer, however, relatively fewer pecks were long enough to qualify for reinforcement, and at a criterion of 48 seconds only a very small proportion of pecks was reinforced (rear distribution).

The tendency for responding to exceed the reinforcement criterion at short latencies and to fall below it at long latencies is a common feature of many performances generated by differential reinforcement of temporal properties of behavior, including performances of humans as well as those of pigeons and rats (Catania, 1970). This may depend in part on the typical asymmetry of such procedures. The procedure places a lower limit but no upper limit on reinforced latencies; latencies well above criterion at short minimum reinforced latencies would presumably be shorter if an upper as well as a lower limit were imposed (e.g., instead of reinforcing any peck with a latency of two seconds or more, we might reinforce only pecks with latencies between two and three seconds).

Schedules of differential reinforcement

The differentiation of temporal properties of behavior is relevant to schedules of reinforcement, treated in detail in chapter 8. We will see that substantial

quantities of behavior can be maintained even when only an occasional response produces a reinforcer. In reinforcement schedules, responses can be made eligible for reinforcement on the basis of the number of responses already emitted, the time elapsed since some event, or some combination of these conditions. A property of particular interest is the separation of responses in time. For example, one will be successful in starting a car with a flooded engine only if one waits long enough after the last attempt to start it; conversely, one will be successful in inflating a bicycle tire with a pump that has a leaky connection to the valve only if one pumps rapidly enough. Analogously, reinforcement might be arranged for a pigeon's key peck only if that peck had been preceded by ten seconds of no pecking, or only if pecking had been emitted so rapidly that at least twenty pecks had occurred within the last five seconds. In the first case, pecking decreases; in the second, it increases. Yet to concentrate solely on the pecks might lead us to say that the first case was not an instance of reinforcement. In fact, the unit of responding reinforced in the first case was not a peck but rather a sequence including a pause plus a peck.. To the extent that this combination increased in frequency (thereby demonstrating that the reinforcers had affected the response class that produced them), the rate of pecking necessarily decreased.

The procedure in the first example, in which a response is reinforced only if preceded by a minimum time without a response, is called a *differential reinforcement of low rate* or *DRL* schedule (Ferster and Skinner, 1957). The procedure is sometimes also called a schedule of interresponse time or *IRT* reinforcement (Malott and Cumming, 1964), because reinforcement is based on the spacing in time of individual responses rather than on the average rate generated by many responses occurring over an extended time. In general, the longer the interresponse time or *IRT* required for reinforcement, the lower the rate of responding (Wilson and Keller, 1953; Staddon, 1965). In *DRL* performance, the rate of responding decreases because the likelihood of responses preceded by long pauses (long *IRT*s) increases. Thus, in *DRL* responding the *IRT* must be regarded as a component of a complex operant consisting of an *IRT* plus a response, in that order.

In the second example, in which reinforcement depended on the emission of twenty or more pecks within no more than five seconds, responding might increase, but only because the high rate of responding itself had been differentially reinforced. Such a schedule is called a *differential reinforcement of high rate* or *DRH* schedule. The *DRH* schedule has received less attention than the *DRL* schedule. The reason is that it is harder to work with *DRH* than *DRL* schedules. Consider the pigeon whose pecking has been raised to a rate of more than four pecks per second by a *DRH* schedule that requires twenty pecks per five seconds. (In fact, by gradually raising the criteria a skillful experimenter can routinely produce rates in excess of ten pecks per second in pigeons.) As long as the pigeon maintains this rate, pecks will produce reinforcers frequently. But if for any reason the pigeon slows down, pecking at a rate that less often meets the reinforcement criterion, pecks will produce reinforcers less frequently. This in turn reduces the pigeon's rate of pecking further, which leads again to a decrease in reinforcer frequency. This vicious cycle may end in a rate of pecking so low that the reinforcement criterion is never met, and the pigeon's pecking may cease completely. To

reinstate responding, the experimenter would have to establish the high-rate performance again through shaping.

In the DRL schedule, however, an effect of reinforcers following pecks is that pecks occur more often. This higher rate of pecking means that pecks occur closer in time, and therefore fewer follow IRTs long enough to meet the reinforcement criterion. Thus, the delivery of reinforcers decreases, and the rate of pecking decreases in turn. But the decrease in rate simply makes IRTs long enough to meet the reinforcement criterion more likely. Thus, pecking oscillates between increased rates accompanied by decreased reinforcement and decreased rates accompanied by increased reinforcement. Pecking is maintained over long periods of time by such contingencies, allowing differential reinforcement gradually to affect IRTs.

The significance of these examples is that we must be cautious about taking rate of responding as a fundamental measure of the effects of reinforcement. Response rate was once regarded as a measure of response strength (cf. reflex strength in chapter 3; Nevin, 1974), but with the recognition that it was simply a property of behavior that could be differentiated like other properties such as force and topography, this view became less tenable. Although DRL responding occurs at a low rate, this responding is easily established and maintained; and although DRH responding occurs at a high rate, this responding is hard to establish and fragile once established.

Various other classes of differential-reinforcement schedules exist, such as the *differential reinforcement of paced responding*, which sets both an upper and a lower limit on the IRTs that can precede a reinforced response and which tends to generate a fairly constant and continuing response rate. The *differential reinforcement of other behavior* or of *zero behavior* (usually abbreviated DRO) delivers a reinforcer if a specified time elapses without a designated response. Consider, for example, a DRO schedule of ten seconds. This schedule reinforces the class of all responses consisting of ten seconds with no pecking. But if a reinforcer is delivered after every ten seconds without a peck, then each peck will delay the reinforcer by at least ten seconds. If pecks then occur less often because they prevent or delay reinforcers, it would be appropriate to say pecks were negatively punished. These ambiguities were considered in chapter 5, in which we considered whether it was worthwhile to distinguish between the positive reinforcement of not-responding and the negative punishment of responding. We concluded there that speaking in terms of discrete responses was less likely to mislead than speaking in terms of the absence of those responses. In name, the DRO schedule seems to raise similar questions, but perhaps it will be less troublesome simply because it involves specifying a procedure rather than attributing functional properties to such events as not-responding.

Complex behavior: Maze learning

So far, we have spoken mostly of relatively simple responses, such as the rat's lever press or the pigeon's key peck. But our account of differentiation indicates that complex sequences of responses can also be treated as operant classes. For

example, a rat's negotiation of a maze might be regarded as a single but very complex response. Consider the simplest maze, the T-maze (Figure 4-2). If the left goalbox is empty and the right goalbox contains food, a response sequence that includes running from startbox to choice point, turning right and then running to the right goalbox will be reinforced. Sequences that include turning left will not be reinforced or (if the rat is allowed to retrace its way from left to right goalbox) will at least not lead to the reinforcer as quickly. As the rat comes more and more frequently to choose the right path we may say that this T-maze performance has become differentiated.

But maze learning can involve more complex sequences. Consider the differentiation of sequences of responses in the *double-alternation problem* (e.g., Hunter, 1928). At one end of a central runway, a choice point provided an opportunity for a left turn (*L*) or a right turn (*R*), but both paths led back to the other end of the runway. Without being removed from the apparatus the organism could come back up the central runway and again choose left or right at the choice point. After a series of right and left turns in the appropriate order, the experimenter could allow the organism to find food as it returned to the central runway. In this apparatus, the question was whether organisms could learn a double alternation sequence, either *LLRR* or *RRLL*. One function of these experiments was to compare the behavior of different species, and in fact the sequence was learned with varying degrees of success by different species and by different organisms of the same species. A more general issue was whether the organism's behavior at one time could determine its behavior at a later time. An organism could not complete a double-alternation sequence successfully unless its behavior at the choice point somehow took into account the turns it had made the preceding times around.

Aside from the length and complexity of the sequence, a major difference between this and the earlier examples of differentiation is in specifying the response class. Linear position or force or duration are single dimensions along which particular responses can be located, but sequences of responses in double-alternation learning cannot be ordered unambiguously along a single dimension (e.g., the sequence *LLLRR* contains only a single transition from *L* to *R* and the sequence *LRLRR* contains as many *L*'s as *R*'s, but then which is more closely related to the reinforced sequence, *LLRR*?). An even more serious question is whether to treat a sequence of turns as a single unit of behavior or to subdivide it into smaller components (e.g., if there are more *R*'s than *L*'s in learning the *LLRR* sequence because only *R* is ever followed immediately by a reinforcer, should individual turns rather than entire sequences be treated as units?). The study of motor skill in humans sometimes deals with the acquisition of similarly complex response sequences (e.g., Fitts, 1964; Bilodeau, 1969).

As an organism learns response sequences in these and related tasks, differential reinforcement acts upon the behavior that the organism brings to the experiment. This behavior is often systematic. For example, a rat learning a maze first might always take left turns, then later always take right turns, then still later alternate between left and right, and so on until various components of the maze

are mastered. In a discrimination task with light and dark alleys, in which light-alley choices led to food (i.e., a maze in which correct choices were correlated with discriminative stimuli), a rat at first might mostly choose the dark alley, then mostly the left alley, and then mostly the right alley, until finally it began consistently choosing only the light alley. These systematic patterns were noted by Krechevsky (1932), who called them hypotheses and discussed their implications as follows:

Almost every description of animal learning includes such phrases as "random exploratory movements"; "chance entrances"; "chance errors." It is either implicitly assumed or explicitly stated in most descriptions of learning that in the beginning of the process the animal is a 'chance' animal. His responses are without purpose, without form, and without meaning. Helter-skelter trial and error seems to be the rule at first, and then, after such behavior has eventually led the animal to experience the "correct" pattern, the various laws of learning step in to stamp in the correct responses and stamp out the incorrect. (Krechevsky, 1932, p. 157)

He then contrasted this view with a different one, based upon his data:

The animal . . . brings to each new situation a whole history of experiences. These experiences the animal is ready to apply. *From the very beginning*, perhaps, the animal goes about solving his problem in a straightforward, comprehensive manner wherein each response is . . . a meaningful part of his total behavior. The animal, in executing a series of movements which we call "perfect," "errorless," "learned," "integrated," is not doing something which has arisen from a series of "imperfect," "unintegrated," "chance" responses. He is now merely running through a different set of integrated responses, which series of integrated responses were *preceded by other just as integrated responses*. Such responses, "false solutions," "early systematic attempts," etc., we have dubbed with the dubious name of "hypotheses." . . . When the human individual *behaves* in the very same such and such way we must also say that he has an "hypothesis." However, we are primarily interested not in defending our terminology but in describing certain behavior. The term "hypothesis" has merely been chosen as a convenient tag for such behavior. (Krechevsky, 1932, pp. 528-29; cf. Levine, 1966)

Response sequences: Chaining versus temporally extended units

Once we break down a behavior sequence into components, we may begin to treat the sequence as a succession of different operants, each defined by the reinforcing consequence of producing an opportunity to engage in the next until the sequence is terminated by a reinforcer. Such a sequence is called a response *chain*. An example was provided, in chapter 4, by Sidman's detailed analysis of a rat's lever pressing reinforced by food. Rising up to the lever produced contact with the lever, which set the occasion for pressing the lever, which produced a seen food pellet, which set the occasion for moving to the food cup, and so on. At each stage,

behavior is consequential, and any given segment serves the dual function of reinforcing the last response and producing the conditions that occasion the next one.

Some behavior sequences reasonably seem to be reducible to smaller units in this way, and an analysis into such components can be confirmed experimentally simply by seeing how independent the components are from each other (cf. Skinner, 1934b). For example, if pressing the lever no longer produces a food pellet, we may find that lever pressing decreases, but we may also find, by delivering food pellets independently of lever pressing, that a seen food pellet continues to occasion movement to the food cup. This procedure demonstrates that the integrity of one component is not affected by altering the reinforcement contingencies for another. We could also find out about components by determining whether substitutions can be made. For example, we might eliminate the seen food pellet by conducting the procedure in darkness and find that the sound of a food pellet dropping into the food cup could take over the former function of the seen food pellet.

The critical question is not so much whether some behavior sequences are held together in this way, but rather whether this is the only basis for the development of behavior sequences. In the history of the psychology of learning, the positions taken on this issue were often symptomatic of serious divisions among researchers of different orientations, some holding that sequential behavior could always be interpreted in terms of the concatenation of such components (variously called associations, chains, S-R or stimulus-response bonds, or conditioned reflexes, among others), and others holding that sequential behavior could not be interpreted adequately in such terms. We will conclude that sequential behavior of both sorts is possible; the significant experimental question is to determine the nature of the serial ordering in any particular case.

Perhaps the most telling argument was made by Lashley (1951), who summarized earlier conceptions as follows:

... the only strictly physiological theory that has been explicitly formulated to account for temporal integration is that which postulates chains of reflexes, in which the performance of each element of the series provides excitation of the next. This conception underlay the "motor theories" of thinking which ... sought to identify thought with inaudible movements of the vocal organs, linked together in associative chains. The peripheral chain theory of language ... described speech as a succession of vocal acts in which the kinesthetic impulses from each movement serve as a unique stimulus for the next in the series. (Lashley, 1951, p. 114)

Lashley then described cases, including illustrations from both language and music, that made such accounts implausible.

For example, in answer to the argument that each movement serves as a unique stimulus for the next, Lashley considered the complex sequence of movements required to pronounce the sounds of the word "right" in proper order. The order is not given by the sounds themselves, because the sounds can occur in a

variety of orders and combinations (e.g., in the opposite order, as in "tire"). Thus, the sound sequence cannot be based solely on direct connections, but must depend on some larger organization. Lashley extended the case from sequences of sounds within words to sequences of words within sentences:

The word "right," for example, is a noun, adjective, adverb, and verb, and has four spellings and at least ten meanings. In such a sentence as "The millwright on my right thinks it right that some conventional rite should symbolize the right of every man to write as he pleases," word arrangement is obviously not due to any direct associations of the word "right" itself with other words, but to meanings which are determined by some broader relations. (Lashley, 1951, pp. 115-16)

And Lashley used music to point out that the sheer rapidity of some sequences constrained how the sequences might have been generated:

The finger strokes of a musician may reach sixteen per second in passages which call for a definite and changing order of successive finger movements. The succession of movements is too quick even for visual reaction time. In rapid sight reading it is impossible to read the individual notes of an arpeggio. The notes must be seen in groups. . . . Sensory control of movement seems to be ruled out in such acts. (Lashley, 1951, p. 123)

Lashley's argument, then, was that some sequential patterns of responding cannot be reduced to a succession of stimulus-response or S-R units. When a skilled typist rapidly types the letters *the*, these letters cannot be discriminative stimuli for the next stroke, first because the typist will be executing that next stroke even before the typed letters on the page can have any stimulus effects, and second because these letters cannot be a unique discriminative stimulus if they can occasion hitting the space bar or various other keys depending on whether the typist is typing the word *the* or *these* or *then* or *thermometer*.

The historical problem was probably that, in the face of such arguments, researchers felt that a choice had been forced between assuming that sequential behavior depended upon stimulus-response sequences and assuming that it depended on temporally extended units of behavior not reducible to such sequences. Instead, we may regard the issue as an experimental one. Clearly some sequences can be put together in such a way that each response produces stimulus conditions that set the occasion for the next response, whereas others must be integrated in such a way that the responses appear in the proper order without each depending on the consequences of the last. For any given behavior sequence the issue is deciding which type of sequence it is.

Limits on differentiation

Some properties of responding are more easily manipulated by differential reinforcement than others. Sometimes the basis for such differences raises no theoretical questions. For example, we do not expect flight to be the same in

pigeons and bats and bees, and we ordinarily do not even ask about the possibility of flight in the rat. The differences are obviously anatomical, and they illustrate how much we take limitations on an organism's motor or sensory capacities for granted. But a difference in the effect of differential reinforcement on different responses is more likely to raise questions when it has no clear anatomical basis.

For example, like motor capacities the capacities of various stimuli to reinforce differ across species. Before the relativity of reinforcement was recognized, findings in which established reinforcers for a species failed to have their characteristic effects were hard to deal with. Breland and Breland (1961) described several cases and used them to argue against the generality of reinforcement as a behavioral process. In one demonstration with raccoons, food reinforcers were delivered when a raccoon picked up coins from the ground and deposited them in a container. After some repetitions of the procedure, the raccoon began persistently to rub the coins together instead of releasing them into the container. Contrary to the Brelands' argument, this finding did not invalidate the principle of reinforcement. A raccoon ordinarily rubs and washes its food before eating it, and the coins apparently provided a better opportunity for this behavior than the food that was supposed to function as a reinforcer. The demonstration showed that rubbing was sufficiently more probable than eating that food was not effective as a reinforcer; and we might also assume that an opportunity for rubbing could have been effective as a reinforcer for other responses, perhaps including eating.

Demonstrations of limitations on properties of behavior that can be differentiated do not invalidate our analysis, because our account has been mainly concerned with naming phenomena and relating them to one another. When an experimental procedure is applied, there is no guarantee that the procedure will be effective. A procedure that leads to learning with one response or one organism may not do so with a different response or a different organism, and a response or an organism affected by one procedure may be unaffected by other procedures. No single process accounts exhaustively for the phenomena of learning, and part of the task of an analysis of learning is to explore the limits on the processes that we study.

SUMMARY AND EXTENSION: THE STRUCTURE OF BEHAVIOR

We have seen how behavior can be changed through shaping and how differential reinforcement is involved in generating response classes. The close correspondence between the class of responses with consequences and the class of responses generated by these consequences is the criterion for speaking of an operant class. These classes may be defined along single dimensions such as force or location, or may have more complex properties. Our examples included differential-reinforcement schedules and integrated response sequences. Questions about differentiable properties of responding are concerned with the structure of behavior. When we

describe such properties, we must include how they are measured; the measuring techniques may vary from one response to another.

In each case, our major interest is with the dimensions along which responding conforms to the class of responses that is reinforced. The structure of behavior is such that we cannot always define such dimensions independently of the reinforcement contingencies. For example, consider the reinforcement of novel or original responses in the porpoise (Pryor, Haag, and O'Reilly, 1969). These investigators produced novel performances by selecting for reinforcement, in each experimental session, some class of responses not reinforced in any previous session. After several sessions, the porpoise began to emit responses, in each new session, that the experimenters had never seen before. Response novelty had been differentiated, but how else is this operant class to be specified except by describing the criteria for reinforcement? The fact that we have difficulty measuring them does not rule out originality or other complex dimensions of behavior as properties that may define an operant class. Even the variability of responding itself can be a basis for differential reinforcement (Schoenfeld, Harris, and Farmer, 1966).

We will return to these issues when we discuss verbal behavior. Sentences can be differentiated into grammatical classes such as active versus passive voice, or declarative versus interrogative. Verbal behavior is consequential behavior, and the finding that some grammatical structures are more easily learned than others is similar in kind to the finding that rats learn single-alternation sequences (*LRLR*) more easily than double-alternation sequences (*LLRR*). In each case, the problem is to identify the dimensions along which responding may come to conform to the class of responses that has consequences. For any particular organism, operant classes may be more easily established along some dimensions than along others. These dimensions may be simple or complex, intensive or structural, and easy to measure or difficult. Perhaps they can sometimes be specified only by verbal description (e.g., the class of all responses not seen or not reinforced on earlier occasions, as in reinforcing the porpoise's novel responses). A concern with these dimensions is a structural concern, and must be distinguished from a functional concern, such as an analysis of the changes in behavior produced by the differential consequences of responses along these dimensions. We discussed the structure and function of behavior in chapter 1. The present remarks are simply a reminder that these two properties of behavior are not incompatible; They complement each other, and our analysis of behavior will be deficient if we ignore either.

Discriminated Operants: Stimulus Control

THE NATURE OF DISCRIMINATED OPERANTS

The problem of attention

Place-learning versus response-learning

Stimulus-control gradients

Generalization gradients

Post-discrimination gradients

Inhibitory gradients

Fading: Stimulus control by successive approximations

The vocabularies of differentiation and discrimination: A summary

DISCRIMINABLE PROPERTIES OF THE ENVIRONMENT

Matching-to-sample

Concept formation

Natural concepts

Preparedness

Learning set

SUMMARY

SEVEN

The words *differentiation* and *induction* applied to responses classes are paralleled by *discrimination* and *generalization* applied to stimulus classes. These words have no obvious etymological ties. *Differentiation*, from the Latin *dis-*, apart, plus *ferre*, to carry, is related through the Indo-European *bher-*, to carry or bear, to such words as *birth*, *transfer*, and *metaphor* (but not *interfere*). *Induction*, from the Latin *in-*, in, plus *ducere*, to lead, is related through the Indo-European *deuk-*, to lead, to such words as *duke*, *educate* and *conduct*. *Discrimination*, from the Latin *dis-*, plus *crimen*, judgment, is related through the Indo-European *skeri-*, to cut or separate, to such words as *crime*, *describe*, and *criterion*. And *generalization*, from the Latin *genere*, to produce or cause, is related through the Indo-European *gen-*, to give birth or beget, to such words as *ingenious*, *kind*, and *nature*.

Responses can be differentially reinforced not only with respect to response dimensions, but also with respect to the dimensions of stimuli in the presence of which responses occur. For example, a rat's lever presses in light are different from its presses in darkness, and reinforcement can be arranged for lever presses in the presence but not the absence of light. Similarly, a pigeon's key pecks in green light are different from its pecks in red light. If only pecks in green are reinforced, we say that reinforcement of key pecks is correlated with green. A response class established by such differential reinforcement with respect to stimulus properties is called a *discriminated operant*.

Discriminated operants are a pervasive aspect of behavior. When driving, we proceed through an intersection if the traffic light is green but not if it is red. When speaking to someone, what we say is affected by such factors as whether the other person is speaking, what the other person says, the listener's posture and facial expression, and the setting of the conversation. Many earlier examples of reinforcement included discriminative control of responding. In discussing the rat's learning of a maze, we emphasized the increase in the choice of correct turns, but the rat that failed to discriminate the appropriate place at which to turn right or left would repeatedly bump the walls as it moved through the maze, and could hardly master the maze as a whole.

Discriminative stimuli correspond to the stimuli colloquially called signals or cues. They do not elicit responses. Rather, they set the occasion on which responses have consequences, and are said to *occasion* or *evoke* responses. An example of the development of stimulus control, or of the control of responding by a discriminative stimulus, is shown in Figure 7-1 (Herrick, Myers, and Korotkin, 1959). Rats pressed a lever in the alternating presence and absence of light from a small lamp. When the light was on, lever presses were occasionally reinforced with food. When the light was off, lever presses were not reinforced. The notation for the stimulus correlated with reinforcement is S^D for *discriminative stimulus* or S^+ for positive stimulus; the notation for the stimulus correlated with nonreinforcement or extinction is S^Δ , also for discriminative stimulus (Δ is delta, the Greek letter d), or S^- for negative stimulus (strictly, S^o would be more appropriate, because nothing is removed in the presence of this stimulus, but S^- has become the more common usage).




Figure 7-1. Rats' lever pressing in the presence and absence of a light. The light (S^D) was correlated with variable-interval reinforcement and its absence (S^Δ) with extinction. The discrimination index is the percent of total responses in the presence of the light (100 times S^D rate divided by sum of S^D and S^Δ rates). The data are median rates from eight rats. (From Herrick, Myers, and Korotkin, 1959, Figure 2)

In the procedure of Figure 7-1, light and dark alternated irregularly; when the lamp was lit, it remained on for varying periods ranging from five to thirty seconds. In addition, lever presses were reinforced according to a *variable-interval* or VI schedule of reinforcement: on the average, only one lever press per thirty seconds was reinforced when the light was on. For the present, the important features of this schedule are that it maintains responding and that the varying times between successive reinforcers make the time since the last reinforcer an unreliable predictor of when the next response can be reinforced. In these circumstances, if the rat presses more in the light than in the dark the light may be regarded as the controlling stimulus; such a rate difference could not depend on the time between stimulus changes or the time between reinforcers because the temporal cycle of these events was unsystematic.

Figure 7-1 shows that over sessions lever pressing in light increased whereas in its absence it decreased. The increase is a frequent accompaniment of such discrimination procedures. The figure also shows changes in a discrimination index, responding in the light as a percentage of total responses. The index increased over sessions. This outcome may be stated in many ways: lever pressing in the presence of light was a discriminated operant; the light functioned as a discriminative stimulus for lever pressing; or, lever pressing was under the stimulus control of the light.

THE NATURE OF DISCRIMINATED OPERANTS

We can illustrate some characteristics of discriminated operants with a hypothetical example related to that of Figure 6-2. Again we have a rat in a chamber with a slot in one wall, but this time the rat cannot poke its nose into the slot. Instead, the slot is covered by translucent plastic, and a series of fifteen lamps behind the plastic can illuminate successive segments of the slot. In other words, the lamps provide a stimulus dimension the components of which are analogous to those of the response dimension of Figure 6-2. Centered beneath the slot is a lever, and the food cup of a pellet dispenser is on the opposite wall. Let us now light the lamps behind the strip for fixed durations but in irregular order and record the lever pressing that occurs in the presence of each. If we do not reinforce lever presses, we can expect pressing to be infrequent, and we would not expect the presses to bear any systematic relation to the portion of the strip that is lit. In fact, we might expect data much like those in *A* of Figure 6-2, with the main difference that the *x*-axis would represent stimulus position instead of response position.

At this point, we begin to reinforce lever presses only if the slit is lit at positions 9, 10, 11 or 12; presses are not reinforced when a light comes on in any other position. The initial effect of reinforcement is much like that shown in *B* of Figure 6-2: an increase in responding to lights at all positions. The effect of reinforcement is not restricted only to stimuli at positions correlated with reinforcement; instead,

the effect of reinforcement at some positions spreads to other positions. This phenomenon, the spread of the effect of reinforcement in the presence of one stimulus to other stimuli not correlated with reinforcement, is called *generalization*. In the example, reinforcing responding in the presence of light at positions 9 through 12 affected responding in the presence not only of those lights but also lights at all other positions across the entire slot. The main difference between this example and the original example is in the dimension correlated with reinforcement: in chapter 6 we dealt with a response dimension whereas we are now dealing with a stimulus dimension.

Suppose we continue differential reinforcement with respect to stimulus location, reinforcing responses in the presence of lights at positions 9 through 12 but not at other positions. Responding gradually increases with lights at positions correlated with reinforcement, and decreases with lights at other positions. Again the effects are like those of Figure 6-2, in *C*, *D* and *E*. Eventually most responses occur with lights at positions correlated with reinforcement, as in *E*, and a point may be reached at which, even though some responses still occur with lights at other positions, the distribution of responses does not change much with continued differential reinforcement.

In this example, the stimuli in the presence of which responses were emitted came to conform closely to the class of stimuli correlated with reinforcement. This process is called *discrimination*, and responding that has come under the control of stimuli in this way is said to be *discriminated*. Differential reinforcement with respect to stimuli established a response class defined in terms of the stimuli in the presence of which responses occur. But what about responses in the presence of stimuli outside the boundaries correlated with reinforcement (e.g., positions 6, 7 and 8 or positions 13, 14 and 15)? According to a strict interpretation, they cannot be counted as members of the discriminated-operant class; they can be spoken of in terms of generalization. Yet the same operation of differential reinforcement generates responding inside and outside these boundaries, and this responding is one continuous distribution.

The resolution is the same as in chapter 6. We must recognize two classes of stimuli: one is the class of stimuli correlated with a reinforcement contingency; the other is the class of stimuli in the presence of which responding occurs. Our interest is not in either class alone, but rather in the correspondence between the classes. The discussion closely parallels the discussion of differentiation and induction in chapter 6. This is fitting, because a stimulus in the presence of which a response occurs in another property of that response, like its force or its duration or its topography. Why then do we speak of the effects of differential reinforcement with respect to response properties in terms of differentiation and induction, whereas we speak of the analogous effects of differential reinforcement with respect to stimulus properties in terms of discrimination and generalization?

One important methodological factor may be the basis for the distinction. When we study differential reinforcement with respect to response properties, we measure responses in different classes, but aside from arranging contingencies there

is not much else we can do about them. If we see a rat about to poke its nose into position 7 along the slot, we cannot appropriately prevent the response even if the rat has responded many more times at that position than at other positions. Suppose, however, that we were working with stimulus properties. We could present lights equally often at each position, or we could present lights at some positions but not others so that the rat would never have an opportunity to press the lever in the presence of some stimuli. Such cases justify saying that the stimuli in discrimination procedures set the occasion for responses: when a class of responses is defined by the presence of a stimulus, responses in this class cannot occur when the stimulus is absent.

Even this methodological distinction has exceptions. For example, consider the differential reinforcement of interresponse times (e.g., the DRL schedule; chapter 6). If a pigeon's key peck is reinforced only after at least five seconds of no pecking, the pigeon may come to space its pecks about five seconds apart. We discussed this behavior in terms of the differentiation of a complex operant consisting of a pause plus a peck. We could instead treat the duration of the pause as a stimulus property and argue that the behavior should be dealt with as discrimination with respect to time elapsed since the last peck. In fact, in this particular instance the vocabularies are interchangeable; such ambiguities are always possible when the properties of responses make it easy to speak of the organism as discriminating some property of its own past behavior. Whether we speak of differentiation and induction or of discrimination and generalization, the underlying operation in each case is differential reinforcement. Both differentiation and discrimination reduce to the correspondence between the dimensions upon which differential reinforcement is based and the dimensions of the resulting behavior. We will nevertheless maintain the established difference between the vocabularies of response properties and stimulus properties, because these vocabularies have an extensive and widely accepted historical foundation.

The problem of attention

In discussing the correspondence between the stimuli with which reinforcement contingencies are correlated and the stimuli to which the organism responds, we spoke in terms of the stimulus dimension selected by the experimenter. But stimuli have varied properties, and there are no guarantees that the organism will respond to just those properties selected. In differentiation, a rat's lever presses may come to have fairly constant form even though only force or location is the basis of differential reinforcement. For example, the rat may reliably press the lever with its left paw even though this property is not critical to whether the response will be reinforced. Similarly, in discrimination a rat might respond on the basis of the intensity of a visual stimulus even though differential reinforcement is based only on its location or shape. (Stimulus properties to which an organism is likely to respond discriminatively are sometimes called *salient*, but the property of

salience is based on the organism's behavior with respect to the stimulus rather than on its physical dimensions.)

The concept of *attention* arises from discriminated responding to the properties of stimuli (cf. Egeth, 1967; Mackintosh, 1965). Consider a pigeon whose key pecks are occasionally reinforced with food. One of two stimulus combinations is presented on the key: a triangle on a red background or a circle on a green background. Each stimulus is presented for three minutes. After three minutes of triangle-on-red, the next key peck in the presence of this stimulus is reinforced; after three minutes of circle-on-green, the stimulus turns off without reinforcement. This arrangement in the presence of triangle-on-red is called a *fixed-interval* or FI schedule of reinforcement; the arrangement in the presence of circle-on-green is extinction. We will examine the fixed-interval schedule in chapter 8. For the present, it is sufficient to note that this schedule usually maintains responding that becomes more and more probable as time passes in the interval, rather than the relatively constant rate of responding ordinarily maintained by the variable-interval schedule. If every peck in the presence of triangle-on-red produced a reinforcer, then the reinforcer deliveries themselves might acquire discriminative functions. With FI reinforcement we do not have to worry about such effects of reinforcer deliveries, because no peck can be reinforced until the interval has ended.

Data from this procedure for two pigeons are shown in Figure 7-2 (G. S. Reynolds, 1961b). The graphs on the left show rates of pecking in the presence of each stimulus compound at the end of eighteen hours of training. Both pigeons emitted more than forty pecks per minute during triangle-on-red, but pecked at relatively low rates during circle-on-green. In a test without reinforcement, each component of the compound stimuli was presented separately, with results shown in the right graphs of Figure 7-2. For Pigeon 105, almost all pecking occurred in the presence of the triangle; red, the background color correlated with reinforcement, occasioned little more responding than either the circle or green, the components previously correlated with extinction. For Pigeon 107, on the other hand, almost all pecking occurred in the presence of red; even though the triangle had been correlated with reinforcement during training, by itself this stimulus occasioned even less pecking than the circle or green. We might therefore say that Pigeon 105 was attending to form and not to color, and that Pigeon 107 was attending to color and not to form. Under the training conditions, both form and color of the stimulus compounds were correlated in the same way with reinforcement. Only by separating the components and examining their individual effects could we say what their discriminative functions were.

The phenomenon here called attention is not primarily an instance of generalization. For Pigeon 105, for example, responding did generalize from triangle-on-red to triangle without a colored background, but it did not generalize to red without a triangle. We do not speak of attention as the organism's response to particular stimuli; instead, we speak of the organism's attention to stimulus dimensions (but cf. chapter 14). We say that Pigeon 104 attended to form and not to color because discriminated responding occurred across changes in form but not across changes in background. When an organism responds discriminatively to one




Figure 7-2. Key pecking of two pigeons during reinforcement correlated with triangle-on-red and extinction correlated with circle-on-green (left: training), and during extinction tests in which the components of these stimuli were presented separately (right: attention tests). The tests show that Pigeon 105 was responding discriminatively to form but not color whereas Pigeon 107 was responding to color but not form. (From G. S. Reynolds, 1961b, Figure 1.)

stimulus dimension but not to another, it is appropriate to say that the organism is attending to the first dimension but not the second.

Once an organism has attended to particular stimulus properties in one situation, it is also likely to attend to those properties in a new situation (e.g., Lawrence, 1949, 1950). We might also expect that we could change the likelihood that an organism would attend to one stimulus property or another simply by changing the way in which reinforcement is correlated with these properties. In the demonstration of Figure 7-2, form and color were correlated equally well with reinforcement. For this reason, the experiment was particularly suitable for a demonstration of some properties of attention. But if form discrimination in the pigeon were our major interest, we could correlate stimuli along this dimension with reinforcement and make color irrelevant to reinforcement.

Place-learning versus response-learning. The preceding issues are indirectly related to a long-standing controversy in psychology, on place-learning versus response-learning (e.g., Restle, 1957). We can differentiate the right turns of a rat at the choice point of a *T*-maze by reinforcing right turns but not left turns. We may then ask whether the rat's responding should be characterized by response dimensions (movements to the right as opposed to movements to the left) or by stimulus dimensions (movements toward a particular place, without regard to the direction from which the rat approaches). For example, suppose the right arm of the *T*-maze pointed toward the windowed east wall of a laboratory. The rat might learn right turns; on the other hand, the rat might learn to run toward the brighter windowed side of the laboratory. We could evaluate these alternatives by turning the *T*-maze around so that the right arm now points west and the rat approaches the choice point from the north. If the rat turns right, and therefore away from the windowed side of the laboratory, we would say that the rat shows response learning. If the rat turns left, toward the windowed side of the laboratory, we would say that the rat shows place learning: it moved toward the same place, even though it did so by turning in a different direction. The question is therefore whether the rat has learned right versus left turns or east versus west turns.

The rat's performance depends in large part on the stimuli available both inside and outside the maze. The typical maze used to be topped by wire mesh or some other cover that allowed the experimenter to watch what the rat was doing. But if the experimenter could look in, the rat also could look out. Even though rats are nearsighted, the typical rats can discriminate the general direction of lights and other gross features of the laboratory. As long as stimuli from outside the maze are available, they might become the basis for the rat's turning in a particular direction. But if these stimuli are eliminated by placing an opaque cover over the maze, the direction that the maze faces in the room becomes irrelevant and the rat cannot do anything but demonstrate response learning. Place learning or response learning therefore depend not on any special features of how organisms learn, but rather on how the experimenter prepares the problem for the rat. By a judicious choice of conditions, an experimenter can make either outcome more likely than the other. The problem is the degree to which the organism attends to stimuli inside the maze, perhaps including some aspects of its own behavior, or to stimuli outside the maze. Although what is learned becomes more complex, there are similar issues in cases in which rats choose the shorter available route through a maze when an earlier preferred route has been blocked, or take an appropriate shortcut through a maze even though that shortcut had never before been available (e.g., Tolman, 1948; Appleyard, 1970).

Stimulus-control gradients

The methodology of discrimination procedures places a heavy burden on the experimenter. In a differentiation procedure the order in which responses occur is determined by the organism, but in a discrimination procedure the experimenter

must decide on the order in which stimuli are presented. If an experimenter is interested in some stimulus continuum (i.e., some dimension along which stimuli can vary, such as intensity or position of a light), it may make a difference whether many or few stimuli are presented, or whether each stimulus is presented for a short or a long time, or whether the stimuli include those correlated with both reinforcement and nonreinforcement or only those correlated with reinforcement. Different arrangements of stimuli are the basis for extensive research on gradients of stimulus control.

Generalization gradients. If a response is reinforced in the presence of a single stimulus and some property of that stimulus is varied, responding will depend on the difference between the original and the current stimulus. For example, if a pigeon's key pecks are reinforced when the key is lit yellow, the pigeon will continue to peck, though ordinarily at lower and lower rates, as the key light is changed to orange and to red and to violet. This demonstrates generalization: The effects of reinforcement in the presence of yellow spread to other colors.

Figure 7-3 presents data on generalization of pigeons' key pecks to tones of different frequency after pecks were reinforced only in the presence of a tone of 1000 cycles per second (H. M. Jenkins and Harrison, 1960). In one procedure (no discrimination training), the tone was always present and pecks were reinforced according to a variable-interval or VI schedule. In a second procedure (presence-absence training), the tone was sometimes present and sometimes absent, and pecks were reinforced according to the VI schedule only in its presence. After training, reinforcement was discontinued and tones of other frequencies were presented for the first time, along with no tone and the original training tone. These stimuli were presented eight times each in mixed order, and the pecking in the presence of each stimulus demonstrated the control that each exerted over responding.

As shown for the three pigeons in Figure 7-3 (top), without discrimination training neither the frequency of the tone nor its presence or absence had any substantial effect on pecking. Responding was evenly distributed across frequencies: the gradient of generalization was relatively flat or, in other words, the effect of reinforcement in the presence of the original tone spread uniformly to all of the stimuli tested. But for the five pigeons given discrimination training, in Figure 7-3 (bottom), the original frequency controlled higher rates of pecking than any other frequency; in general, the greater the difference between the original frequency and the tested frequency, the lower the rate of pecking. It is not surprising that low rates of pecking occurred in the absence of the tone, because this was the stimulus condition correlated with extinction. But response rate varied with tone frequency even though discriminated responding had been established only with respect to the presence or absence of the tone and not with respect to frequency. Both types of gradients in Figure 7-3 can be called generalization gradients, because neither followed discrimination training between stimuli on the dimension along which the gradient was obtained. (Differences in the form of generalization gradients are also produced by other variables, such as the reinforcement schedule used during training: e.g., Hearst, Koresko, and Poppen, 1964; D. R. Thomas and Switalski, 1966; or the level of deprivation: e.g., W. O. Jenkins, Pascal, and Walker, 1958).




Figure 7-3. Generalization gradients for individual pigeons as a function of frequency of a tone after the reinforcement of key pecks in the presence of a 1000 cycle per second tone (top: no discrimination training), or after reinforcement in the presence of that tone and extinction in its absence (bottom: presence-absence training). Gradients were relatively flat without discrimination training but were relatively peaked with a maximum at the prior reinforcement stimulus (S^D) after presence-absence training. (From H. M. Jenkins and Harrison, 1960, Figures 1 and 2.)

The concept of attention has been invoked to treat such differences in gradients. The argument is that the tone versus no-tone training produced a peaked generalization gradient to tone frequency because it caused the pigeon to attend to auditory stimuli. There is nothing wrong with saying that the pigeon attends to the frequency of tones with but not without presence-absence training; attention is a name for this kind of phenomenon. The problem is when the term suddenly appears as an explanation. Attending to stimuli is undoubtedly something that

organisms do. We decide whether it is appropriate to say that an organism is attending to a stimulus dimension on the basis of the data, and we should therefore not then use this term, derived from the data, as an explanation for the data that justified it in the first place. (Similarly, if we say that stimulus dimension *A* is more salient than stimulus dimension *B* because organisms are more likely to attend to *A* than to *B*, we cannot then use this salience as an explanation of why the organism attended to *A* rather than to *B*.)

Post-discrimination gradients. Stimulus-control gradients can also be obtained after establishing a discrimination between two or more stimuli along the dimension. Figure 7-4 compares such a post-discrimination gradient with a generalization gradient (Hanson, 1959). For one group of pigeons (generalization), key pecks in the presence of a single wavelength on the key were reinforced according to a VI schedule; the rate of pecking in the presence of a variety of different wavelengths was then determined during nonreinforcement. For a second group (post-discrimination), key pecks were reinforced according to a VI schedule in the presence of the same wavelength as the first group, but this wavelength alternated with another in the presence of which pecks were not reinforced (extinction); as in the first group, the rate of pecking in the presence of a variety of wavelengths was then determined during nonreinforcement. The gradient after training with a single stimulus had its peak at that stimulus; rate of pecking decreased with increases in the difference between the training stimulus and the test stimulus. After discrimination between two stimuli, however, the peak of the gradient was displaced (a peak shift) from the reinforcement stimulus in a direction away from the extinction stimulus. Relative to the gradient obtained from the first group, stimuli closer to the reinforcement stimulus produced higher response rates, and stimuli closer to the extinction stimulus produced lower response rates. (Similar effects are also obtained when the discrimination is based not on reinforcement in the presence of one stimulus and extinction in the presence of another, but rather on a higher frequency of reinforcement in the presence of one than another stimulus: Guttman, 1959).

One account of the form of the stimulus-control gradient after discrimination training is provided by Spence (1937). Spence assumed that reinforcement in the presence of one stimulus created a gradient of increments in responding centered on the reinforcement stimulus (excitatory gradient), that extinction in the presence of the other stimulus produced a gradient of reductions in responding centered on the extinction stimulus (inhibitory gradient), and that the effects of different stimuli after discrimination training could be derived by subtracting the inhibitory from the excitatory gradient. Spence's theoretical difference gradient showed a peak shift: The gradient maximum was displaced from the reinforcement stimulus in a direction away from the extinction stimulus. But subtraction could not make the difference gradient anywhere higher than the original excitatory gradient. This account was consistent with the form but not the absolute value of the post-discrimination gradient. Nevertheless, Spence's account was the foundation for subsequent quantitative analyses of stimulus-control gradients (cf. Hearst, Besley, and Farthing, 1970; Rilling, 1977).




Figure 7-4. Generalization and post-discrimination gradients. The generalization gradient shows pigeons' key pecking in the presence of different wavelengths after reinforcement in the presence of a wavelength of 550 millimicrons (S^D). The post-discrimination gradient shows key pecking in the presence of different wavelengths after reinforcement in the presence of 550 millimicrons (S^D) and extinction in the presence of 570 millimicrons (S^Δ). (The spectrum goes from violet and blue at the short wavelengths to red at the long wavelengths; the training stimuli were in the green-yellow region of the spectrum.) The generalization gradient shows a peak maximum responding is shifted to a wavelength different from the original reinforcement stimulus, in a direction away from the extinction stimulus. (From Hanson, 1959, Figure 1.)

The recognition that stimulus-control gradients could be altered by discrimination training led also to questions about the origins of peaked generalization gradients. One suggestion was that the steepness of a generalization gradient depended on the discriminability of stimuli at various points along the dimension. For example, in the generalization gradient of Figure 7-4, response rate decreases more rapidly at wavelengths below than above the reinforcement stimulus (e.g., the lower rate at 520 than at 580 millimicrons). But when the width and steepness of generalization gradients around stimuli in different regions of the spectrum were compared with thresholds for changes in wavelength in these different regions, no simple relations between generalization and discriminability were demonstrable (Guttman and Kalish, 1956).

Another suggestion was that the gradients were simply the result of discrimination learning before the organism was brought into the experimental situation. For example, a pigeon may learn discriminations among the different colors of the grain that it eats long before it sees yellows or reds or greens in a generalization test. Thus, the natural environment might favor the learning of some discriminations over others. If yellow predominates among the colors of the grain that the

pigeon eats, the pigeon's discrimination might be sharper in the yellow than in other regions of the spectrum even before any experimental conditions. Because we cannot know what color discriminations might be established in an organism's natural environment, an alternative approach is to create an environment in which color discriminations are not possible. This can be done by raising an organism in a monochromatic environment, i.e., an environment lit by only a very narrow band of wavelengths (e.g., the yellow light emitted by a sodium vapor lamp). To humans in such an environment, objects have no color; everything appears in shades of gray.

A suitable organism for an experiment in monochromatic rearing is the duckling, because the duckling is capable of walking, pecking, eating, and drinking soon after it hatches, and reinforcement procedures therefore can be initiated early. Flat generalization gradients across wavelengths were obtained from ducklings that had been reared monochromatically, suggesting that color experience is a condition for peaked generalization gradients for wavelength (N. Peterson, 1962). This finding, however, has not been repeated reliably by other investigators (e.g., Tracy, 1970). In addition, discrimination between only two wavelengths after monochromatic rearing produced fairly standard post-discrimination gradients, including a peak shift, even though ducklings under such conditions were mostly responding to wavelengths they had never seen before (Terrace, 1975). The questions therefore remain to some extent unanswered. But once attention to stimuli along some continuum has been established by differential reinforcement with respect to two stimuli along the continuum, we should not be surprised by differences in responding in the presence of other stimuli along the same continuum (cf. Guttman, 1956). Such differences are properties of the stimulus classes or discriminated operants established by differential reinforcement, and such classes are undoubtedly determined by a variety of factors, including the organism's sensory apparatus and its past history of stimulus control.

Inhibitory gradients. Once Spence's account had been invoked to handle the form of post-discrimination gradients, interest arose in finding a direct way to measure the inhibitory gradient directly. The concept of inhibition was controversial, but the difficulties were methodological as well as theoretical. To determine whether a stimulus reduced responding, there had to be some responding in the first place (cf. Honig, 1961). But once the stimulus in the first place had generated responding, whether in the presence of the stimulus or in the presence of many stimuli along the continuum, the effects of reinforcement had generated responding, whether in the presence of a single stimulus or in the presence of many stimuli along the continuum, the effects of extinction would vary considerably with reinforcement schedule, number of reinforcement sessions, and other factors. A procedure was needed that would make the dimensions along which an extinction gradient was determined independent of the dimensions correlated with reinforcement (H. M. Jenkins and Harrison, 1962; H. M. Jenkins, 1965; Hearst, Besley, and Farthing, 1970). Such a procedure is illustrated in Figure 7-5 (Honig, Boneau, Burstein, and Pennypacker, 1963). These investigators worked with one group of pigeons for which the reinforcement stimulus was a vertical line on the key and the extinction stimulus was a lit key without a line, and a second group for which these stimuli were




Figure 7-5. Excitatory and inhibitory stimulus-control gradients after variable-interval reinforcement in the presence of one stimulus (S^D) and extinction in the presence of a second stimulus (S^Δ). Data from pigeons are shown for two different studies (filled and unfilled symbols). In each study, a vertical line was correlated with reinforcement and the absence of a line with extinction for one group, and these stimuli were reversed for the other group. Responding in the presence of various line orientations and in the absence of a line was determined during non-reinforcement. (From Honig, Boneau, Burstein, and Pennypacker, 1963, Figure 1.)

reversed. For both groups, variable-interval (VI) schedules were used during training, and stimulus-control gradients along the dimension of line orientation were obtained during nonreinforcement.

Figure 7-5 shows data for two such groups from each of two studies. When the vertical line was correlated with reinforcement and the absence of a line with extinction, the rate of pecking was lower with greater deviations from vertical (triangles); for these pigeons, the effect is analogous to that of Figure 7-3 (bottom). Just as peaked gradients with respect to frequency of a tone were obtained after discrimination training with respect to presence or absence of a tone, peaked gradients with respect to line orientation were obtained after discrimination training with respect to presence or absence of a line. But when the absence of the line was correlated with reinforcement and the vertical line was correlated with extinction, the rate of pecking was higher with greater deviations from vertical

(circles). In a post-discrimination gradient, the effects of a stimulus are necessarily determined by the location of both the reinforcement stimulus and the extinction stimulus (e.g., if stimulus *B* on a continuum is between reinforcement stimulus *A* and extinction stimulus *C*, it cannot be moved toward one without also being moved away from the other). In the present instance, changing the orientation of the line affected the distance of the stimulus from the vertical extinction stimulus while presumably having no effect on its distance from the reinforcement stimulus, which was the absence of a line.

Questions can always be raised about the dimensions of the stimuli to which the organism is attending. For example, suppose it was argued that the pigeon was looking only at the very upper edge of the key. If the line is shifted from vertical, its top moves to the right or left, and at its upper edge the key more closely resembles a key without a line. It may also be important that organisms are more likely to attend to properties of stimuli correlated with reinforcement (feature-positive discrimination) than to those correlated with extinction (H. M. Jenkins and Sainsbury, 1970). Such stimulus factors could influence the form of a gradient. Our judgment must be based on the pigeon's behavior, and if we wished to evaluate these possibilities, we would have to do so experimentally (e.g., by removing the lower or middle portions of the line in these various conditions and seeing whether that changed the pigeon's behavior: cf. Touchette, 1969).

Inhibitory stimulus-control gradients seem to imply that extinction represents some kind of active suppression of responding rather than simply the passive outcome of discontinuing reinforcement. This view follows from the equivalence between the concentration of the effects of reinforcement in the region of the reinforcement stimulus and the concentration of the effects of extinction in the region of the extinction stimulus. Yet it is difficult to say whether the extinction procedure itself, the stimulus, or some other event should be identified as the inhibitor: accounts in terms of inhibition are not particularly useful if responding is said to be inhibited but no inhibiting event can be specified. The language of inhibition might more appropriately be reserved for cases in which both the inhibited and inhibiting events can be specified (as when increases in the reinforcement of one response reduce some other response, e.g., Catania, Silverman, and Stubbs, 1974).

Fading: Stimulus control by successive approximations

Just as the response properties that define an operant class can be gradually changed by shaping procedures, the stimulus properties that define a discriminated operant class can be gradually changed by analogous procedures called *fading*. Instead of exposing an organism directly to difficult-to-discriminate stimuli, training can begin with easy-to-discriminate stimuli and then gradually move to the more difficult stimuli. For example, for a pigeon's key pecking reinforced with food, it is usually harder to establish a discrimination between vertical and

horizontal lines than between red and green. Once a discrimination between red and green is established, however, the discrimination between vertical and horizontal can be successively approximated by superimposing vertical on red and horizontal on green and then gradually fading out the colors (Terrace, 1963b).

Establishing stimulus control through fading is often effective (e.g., Hively, 1962; Sidman and Rosenberger, 1967; Sidman and Stoddard, 1967), but as with shaping no simple rules exist to determine how rapidly stimuli should be faded in or out in particular situations. For example, after we had superimposed vertical on red and horizontal on green and had partially faded out the colors, we might find that we could remove them completely, leaving responding under the stimulus control of vertical and horizontal. On the other hand, we might instead find that even after the colors were very dim they remained the stimuli to which the organism was attending, so that discriminated responding disappeared each time we dimmed the colors below certain threshold levels. The former would represent successful fading, but the latter would represent failure to transfer the discrimination from color to line orientation.

Just as shaping requires some behavior available to be shaped, fading requires discriminative responding available to be shifted to a new stimulus dimension. Consider, for example, a phenomenon called errorless discrimination learning (Terrace, 1963a). Shortly after pigeons' pecks on a red key had been shaped with food reinforcement, reinforcement in the presence of red was continued according to a variable-interval schedule. Three-minute periods of red alternated with another stimulus, during which pecks could not be reinforced. At first, this other stimulus consisted of a dark key that lasted for five seconds. Gradually, over the course of three sessions, the duration of this other stimulus was increased and it was changed from dark to dim green and then through various intensities of green to a level that to the human eye matched the brightness of the red key. By the end of these conditions, three minutes of red correlated with reinforcement were alternating with three minutes of green correlated with extinction. With each pigeon for which this procedure was used, stimulus control was established with fewer than ten responses to the extinction key in the entire course of training; pecking occurred almost without exception on the red key rather than on the green key. For pigeons in other groups, for which this procedure was introduced later after initial training or for which green was abruptly introduced at full duration and intensity, hundreds and even thousands of pecks occurred on the green extinction key during equivalent periods of training.

The gradual fading in of green was in part effective because turning the red key dark early in training stopped the pigeon's pecking for a few seconds (any abrupt stimulus change might have had such an effect). Pecking was not likely to start again before the five-second extinction stimulus ended. Thus, a difference in responding to the two stimuli was established at the outset, and the gradual changes in the duration and intensity of the extinction stimulus built upon this difference. Once such discriminative responding is established, reinforcement could be made available in the presence of the faded-in stimulus and the organism would never discover it, because it would not respond in the presence of that stimulus. The sense

in which such a performance is errorless is not obvious. We must be cautious about the language of errors in this type of situation; the term *error* implies a judgment about the value of responding that does not seem appropriate to a behavioral analysis. Nevertheless, the fading procedure is not solely of theoretical interest; because of its potential for applications to education and other areas, it may have much practical significance.

*The vocabularies of
differentiation and
discrimination: A summary*

Both differentiation and discrimination involve differential reinforcement. The major difference between these two processes is whether differential reinforcement is imposed on properties of responding or on properties of the stimuli in the presence of which responding occurs. The main implication of this difference is procedural: in studies of differentiation, the experimenter must wait for the organism's responses, whereas in studies of discrimination the experimenter can control the order and duration of stimulus presentations. These and other variables determine the form of stimulus-control gradients, which include generalization gradients, post-discrimination gradients, and inhibitory gradients. Ambiguous cases of differentiation and discrimination can be identified (e.g., should we say that in temporally spaced responding the organism discriminates the passage of time or that the temporal properties of the organism's responses have been differentiated?), but the procedural difference seems to justify the maintenance of this distinction in vocabulary. The vocabulary is summarized in the following:

Differential reinforcement (Operation)	Concentration of effects of reinforcement (Process)	Spread of effects of reinforcement (Process)	Differential reinforcement by approximations (Operation)
With respect to response properties	Differentiation	Induction	Shaping
With respect to stimulus properties	Discrimination	Generalization	Fading

As with other behavioral terms we have considered, these words name certain properties of behavior and of behavioral procedures; they do not function as explanations.

**DISCRIMINABLE PROPERTIES
OF THE ENVIRONMENT**

Discriminated operants are behavior classes defined by the stimuli that occasion responding. Such classes are often identified in our everyday vocabulary, as when we speak of stopping at a red traffic light or answering a telephone. The red light

can vary in brightness and size and the telephone ring can vary in loudness and timbre, but responding is reasonably independent of variations along such dimensions and so we speak in terms of these classes of events rather than in terms of particular instances.

Relatively simple dimensions of stimuli, such as intensity or location, can be the basis for differential reinforcement. The experimental question is whether responding conforms to the differential consequences, in that more responding occurs in the presence of a stimulus correlated with reinforcement than in the presence of a stimulus correlated with nonreinforcement. Differential reinforcement may also be arranged for complex properties of stimuli that are not easily quantified. For example, children learning to read must be able to name the letters of the alphabet. But the properties important for distinguishing among some letters are different from those important in distinguishing among others (e.g., straight line versus curve is important in distinguishing *U* and *V*, but not *V* and *A* or *V* and *N*), and different distinctions are important for lowercase than for uppercase letters (e.g., no pair of uppercase letters has the up-down or left-right reversals that must be mastered to read *b*, *p*, *d*, and *q*). The way in which a child learns to distinguish among letters of the alphabet depends on relations among such stimulus properties as symmetry, curvature, and closure. Those properties essential to discriminating among different letters are called *critical features* (e.g., Gibson, 1965). It is not sufficient, however, to enumerate critical features. For some letters, the upper-case and lower-case forms differ more from each other than they differ from other letters (e.g., *e*, *E*, and *F*, or *h*, *n*, and *N*). Given the multitude of forms, what then defines the class of stimuli that occasions our saying *a* or *e* or *r*? This question is about stimulus structure in learning to read the letters of the alphabet; it is important to remember that questions of stimulus structure are different from questions of stimulus function.

In this text, we often treat stimuli as if they were restricted to concrete objects or environmental events. But although we can learn to respond in consistent ways to objects or events in our environment, we also can discriminate features, sometimes called abstract or relational, that are independent of any particular object or event. The term *stimulus* may therefore function in a more general way, in the sense of some property of environmental events. For example, we might say that a chair had been placed to the right of a table. Although the chair and the table are concrete objects, being-to-the-right-of is not, and yet we can clearly discriminate this relation from being-to-the-left-of. In some discrimination experiments, therefore, relations among stimuli have been the dimensions of interest.

Matching-to-sample

A procedure called matching-to-sample is illustrated as it might be arranged in a three-key pigeon chamber in Figure 7-6 (cf. Skinner, 1950; Ferster, 1960; Cumming and Berryman, 1965; Wright and Cumming, 1971). During an intertrial interval, all the keys are dark. The trial begins with illumination of the center key, the *sample*. Typically a peck is then required on the center key; it is assumed that




Figure 7-6. Diagram of one trial in a matching-to-sample procedure in a three-key pigeon chamber. After an intertrial interval of t seconds, a sample stimulus (green: G) appears on the center key, sometimes called an observing response, turns on comparison stimuli on the two side-keys. One comparison stimulus matches the sample; the other (red: R) does not. If the pigeon pecks a stimulus matches the sample; the other (red: R) does not. If the pigeon pecks the matching comparison stimulus, food is delivered and the intertrial interval of a new trial begins; if the pigeon pecks the nonmatching comparison stimulus, the sample stimulus and the location of the matching comparison stimulus vary from trial to trial.

this peck guarantees that the pigeon looks at the sample stimulus.¹ This response turns on *comparison* stimuli on the two side-keys. One matches the sample stimulus and the other does not. A peck on the matching side-key then produces a reinforcer followed by a new intertrial interval, whereas a peck on the nonmatching side-key is followed directly by the intertrial interval without a reinforcer (sometimes a nonmatching peck also extends the duration of the intertrial interval, on the assumption that such a consequence would punish nonmatching pecks).

The sample stimulus ordinarily changes from trial to trial; for example, it might be green on some trials and red on others. The position of the matching side-key also changes from trial to trial. A frequent component of matching-to-sample

¹This peck is sometimes called an *observing response*, but that term is also applied to cases in which a response produces some discriminative stimulus, as when reinforcement and extinction alternate irregularly for pecks on one key and pecks on a second key briefly produce stimuli correlated with these contingencies: e.g., Kelleher, 1958b; Kelleher, Riddle, and Cook, 1962; D'Amato and Fazzaro, 1966.

is the correction procedure, which arranges that when a trial ends with a peck on a nonmatching side-key the same sample and comparison stimuli are repeated on the next trial. This procedure prevents the development of consistent responding to only one side-key or only one color. For example, suppose a pigeon always pecked the comparison stimulus on the left side-key. If the position of the matching side-key nonetheless alternated irregularly between left and right, these pecks would be reinforced on half the trials, i.e., those with the matching stimulus on the left. Reinforcement on half the trials could be enough to maintain left-key pecking indefinitely. On the other hand, with a correction procedure the pigeon would have to switch to the right key sooner or later, because once a matching stimulus appeared on the right that combination of sample and comparison stimuli would not change until a right-key peck had occurred and been reinforced. Similarly, if the pigeon always pecked red side-keys, these pecks would be reinforced on half the trials unless a correction procedure forced the pigeon eventually to switch sometimes to pecking green side-keys.

Suppose now a pigeon was responding accurately in a matching-to-sample procedure using red (*R*) and green (*G*) stimuli. How should this performance be characterized? Has the pigeon learned to peck red side-keys in the presence of red center-keys and green side-keys in the presence of green center-keys? Or has the pigeon learned to peck the left key in the presence of the configurations *RRG* and *GGR* and the right key in the presence of the configurations *GRR* and *RGG*? Or might we wish to say simply that the pigeon has learned the concept of matching? We might present blue or yellow as a new sample stimulus. If the pigeon matched these colors we would feel more confident about speaking of a concept of matching (in fact, matching in pigeons does not transfer easily to new colors, although the likelihood of such transfer can be modified by the training procedures). Even if we saw matching with new colors, what would we say if we failed to get matching with geometric stimuli such as squares and circles as the sample and comparison stimuli? Would we say the pigeon had learned the concept of color matching but not form matching? If we did so, we at least would have to note that the human concept of matching seems not so limited by the specific dimensions of the matched stimuli.

There are many variations on matching-to-sample. If reinforcement is arranged for pecks on the nonmatching rather than the matching side-key, the procedure becomes an instance of *odddity* responding, because this key is necessarily the odd key of the three keys (in other versions of oddity procedures, no sample is used; with pigeons in a three-key chamber, for example, each trial consists of lighting all three keys, with the odd key any one of them). The relations among stimuli can also be arbitrary, when the procedure is sometimes called *symbolic matching*. For example, we could train the pigeon to peck a green key given a square as a sample and a red key given a circle as a sample. We might then ask whether the pigeon might peck a square given a green sample and a circle given a red sample. (We often expect such reversibility when we deal with words and objects, as when a child who has learned to point to a picture of a car on seeing the word *car* can also point to the word on seeing the picture of a car. This reversibility, however, is not to be taken for granted; cf. chapter 10.) Each example involves a

complex of stimulus relations; such complexes are sometimes called *concepts* (e.g., Kelleher, 1958a).

Concept formation

Concepts have been spoken of as generalization within a class of stimuli and discrimination between classes of stimuli (F. S. Keller and Schoenfeld, 1950). Thus, our concept of red must involve generalization among all stimuli we call red, and discrimination between these stimuli and others we do not call red. Concepts are therefore to classes of stimuli what operants are to classes of responses. (Some theorists speak of responding on the basis of some single property of stimuli as *abstraction*, and restrict concepts to responding on the basis of some combination of properties. We will not be concerned with such a distinction, in part because of its ambiguity. For example, being-to-the-left-of might be regarded either as a single relational property or as a combination of properties necessarily including both a reference point and a stimulus to the left of that reference point.) Whether we are dealing with size, color, shape, particular objects, or any other dimensions of the environment, we must always consider these as classes of stimuli, and therefore concept formation seems coextensive with discrimination learning. The literature on such learning is substantial (e.g., Blough, 1975; Heinemann and Avin, 1973; Mackintosh, 1977; Wildemann and Holland, 1972).

Natural concepts

We already noted that it is often difficult to define discriminative stimuli by physical dimensions. For example, the properties that define the letter *A* vary according to whether the letter is upper case or lower case, and whether it appears as type or as script. The capacity to discriminate among such stimuli exists in animals as well as humans (e.g., discriminations between various forms of the letter *A* and the digit *2* have been established in pigeons: Morgan, Fitch, Holman, and Lea, 1976). But the difficulties of defining stimuli by measurable physical properties are not limited to arbitrary classes established by humans, such as letters and numbers. They exist also with everyday objects and events. What distinguishes dogs from other animals? On what basis do we generalize between a chihuahua and a husky by calling both dogs, while we discriminate between huskies and wolves or foxes? (In quantitative treatments, such classes are called fuzzy sets: Rosch, 1973.)

Procedures have been established in which pigeons learn to discriminate between pictures that contain a human form and those that do not (e.g., Herrnstein and Loveland, 1964). Such discriminations have been called *natural concepts*. In one study (Herrnstein, Loveland, and Cable, 1976), slides were presented on a screen next to a pigeon's key and key pecks were reinforced in the presence of some slides but not others. Some pigeons learned discriminations between pictures with and without trees; others learned discriminations between pictures with and without water; still others learned discriminations between pictures with and without particular people. After training with one set of slides

(e.g., slides with and without trees), the pigeons discriminated among slides that had not been presented before. New slides were sometimes more accurately discriminated than slides used in training. The implications were drawn that:

... we cannot begin to draw up a list of common elements. To recognize a tree, the pigeons did not require that it be green, leafy, vertical, woody, branching, and so on (overlooking the problem of common elements nested within terms like leafy, vertical, woody, and so on). Moreover, to be recognizable as a nontree, a picture did not have to omit greenness, woodiness, branchiness, verticality, and so on. Neither could we identify common elements in the other two experiments. If not common elements, what? No other theory is so easily characterized, though in crude terms an alternative suggests itself. Pigeons respond to clusters of features more or less isomorphic with the clusters we respond to ourselves. The green should be on the leaves, if either green or leaves are present. However, neither is necessary or sufficient. The vertical or branching parts should be the woody parts, although neither of these features is necessary or sufficient either. What we see as trees comprises a complex list of probabilistic conjunctions and disjunctions, the discovery of which would require far more effort than seems justified by any possible benefit. (Herrnstein, Loveland, and Cable, 1976, pp. 298-299)

Preparedness

An implication of the many examples with responses reinforced only by food has been that the particular consequences of responding do not affect the acquisition of discriminations. Yet just as the effectiveness of a reinforcer may depend upon the reinforced response (cf. chapter 4), the effectiveness of discriminative stimuli may depend on the responses and reinforcers in the discrimination procedure. The concept of *preparedness* is based on the observation that discrimination learning may vary with the context of responses and reinforcers within which the learning occurs; thus, organisms are said to be prepared to learn certain discriminations in some contexts but not others (cf. Seligman, 1970; Seligman and Hager, 1972; Schwartz, 1974). The most striking experimental demonstrations have used punishment rather than reinforcement, and therefore discriminations based upon differential punishment illustrate the point (but see Rozin and Kalat, 1971).

Experiments on food aversion show that a rat will stop eating a particular food if the rat becomes sick after eating it (Revusky and Garcia, 1970). Nausea or other systemic consequences of eating can punish eating even though these consequences follow eating only after a considerable delay. In an experiment on the role of discriminative stimuli (Garcia and Koelling, 1966), thirsty rats were allowed to drink water sweetened by saccharin; their drinking was accompanied by clicks and flashes of light triggered electronically by their licks. In other words, the rats drank water that was not only sweet but also bright and noisy. In one group of rats, drinking was followed by levels of X-irradiation that later produced sickness. In a second group, drinking was followed by electric shock. Later, the drinking of each group

was measured when a sweet solution was presented and when drinking was accompanied by noise and light. The results are summarized in Figure 7-7. Rats that had been X-irradiated drank less sweetened water but their drinking was not affected by noise and light. Rats that had been shocked drank less when drinking was accompanied by noise and light but their drinking was not affected by whether the water was sweetened. In other words, when the aversive stimulus was the delayed systemic consequence of X-irradiation the rats learned only its relation to the taste of the water consumed earlier, but when the aversive stimulus was electric shock the rats learned only its relation to the external noise and light that preceded it. The delayed effects of X-irradiation punished the drinking of sweetened water, and the immediate effects of shock punished drinking accompanied by noise and light. Rats are therefore predisposed to learn different relations among discriminative stimuli and contingencies in different situations. It is not enough to say that organisms learn some responses or some stimuli more easily than others; an adequate account must deal with relations among responses and stimuli.

Revusky and Garcia (1970) discuss the implications of these findings in the context of a thought experiment in which the reader finds one hundred dollars:

The \$100 dollars was left by an insane billionaire experimenter because, two hours ago at lunch, you ate gooseberry pie for dessert instead of your usual apple pie. The experimenter wanted to increase the future probability that you would eat gooseberry pie. It is very unlikely that this experiment will be successful. . . . hundreds of events are bound to occur during the two hours between consumption of the gooseberry pie and the receipt of the \$100. The odds are very great that you would have associated one of these intervening




Figure 7-7. Percent of normal fluid intake by rats whose drinking of sweetened water was accompanied by clicks and lights. In one group, drinking was followed by X-irradiation; in a second group, drinking was followed by electric shock. The nature of the punisher that became effective in reducing drinking depended on the nature of the punisher. (From Garcia and Koelling, 1966, as presented by Revusky and Garcia, 1970, Figure 9.)




Figure 7-8. Learning set or learning-to-learn. Eight monkeys learned a variety of two-choice discrimination problems. Each set of points shows the percent of correct responses over the first six trials of each problem for blocks of successive problems. (Note that the y-axis starts at 50%, or chance level.) The rate of learning increased over successive problems. By the last block of problems, responding was almost 100% accurate by the second trial with each new problem. (Adapted from Harlow, 1949, Figure 2.)

events with the \$100. . . . The results of our thought experiment are really shocking. We have selected an arbitrary reward of great potency and a response of which you probably were well aware. . . . [and] we reached the conclusion that the nature of the environment precluded association over a two-hour delay. (Revusky and Garcia, 1970, p. 20)

These investigators then suggest an alternative thought experiment: instead of finding money, the reader becomes sick.

Since the gooseberry pie was new to you, you would probably conclude that the pie caused the illness. Here association over a two-hour delay agrees with our daily experience. . . . But why has a change from a consequence of \$100 dollars to a consequence of sickness changed the situation? . . . There seems to be only one reasonable answer. . . . The fact that infra-humans also can associate over long delays strongly suggests that there is an innate selective association of flavors with physiological aftereffects and, what is more important, a selective failure to associate irrelevant stimuli with toxicosis. (Revusky and Garcia, 1970, p. 21)

In other words, the kinds of consequences that are differentially correlated with stimuli may influence the kinds of discriminations that can be learned.

Learning set

The organism can come to us prepared to learn different discriminations in different contexts, but we must also recognize that the learning of a new discrimination can be modified by the organism's history of discrimination learning. Relations among stimulus properties may come to control responding independently of specific stimuli. A phenomenon called *learning set* illustrates this (Harlow, 1949). Two different objects are presented to a food-deprived monkey in successive trials. Their positions vary from trial to trial. Food is under only one of the objects, and therefore the response of picking it up is differentially reinforced. After a number of trials, the monkey masters the discrimination. Then a new pair of objects is presented, with differential reinforcement again arranged for picking up only one of them. Again a discrimination is established. Another new pair is then presented, and later still another, and so on.

Data averaged across eight monkeys are shown in Figure 7-8. Each set of points shows the percent of correct responses over the first six trials across blocks of successive discriminations or problems. (Note that the *y*-axis begins at 50%, which is chance level.) Over the first eight problems (1-8), correct responses increased gradually; by the sixth trial, they had not even reached 80%. In the next block of eight problems (9-16), correct responses increased more rapidly over trials. Over successive blocks, correct responding rose more and more rapidly over trials, until by the last block of trials (289-344) responding nearly reached 100% accuracy on the second trial of each new problem. In other words, the more problems the monkey had mastered, the more rapidly it mastered each new problem.

In this procedure, the discriminated operant cannot be described simply in terms of the components of a given stimulus pair. When successive discriminations became established so rapidly that monkeys consistently choose the stimulus correlated with reinforcement after a single trial with a new pair, we must characterize performance by relations between stimuli and their correlated consequences over successive problems rather than within single problems. When the early problems are presented, the monkey must learn not only the particular stimulus correlated with reinforcement within each problem, but also many other aspects of the procedure (e.g., that food is correlated with one of the two stimuli rather than with position or some other feature of the setting; that this correlation does not change within problems; that reaching simultaneously for both stimuli will not be reinforced; and so on). During the early blocks of problems we must assume that learning takes place slowly because the monkey is learning many things; only later, when these other aspects of this situation are learned, does the monkey have to learn only which of the two stimuli is correlated with food in a given problem. At this point, we may define the discriminated operant as follows: If the response to a given stimulus is reinforced on the first trial with a new problem, responding to that stimulus will continue on all subsequent trials; but if that response on the first trial is not reinforced, then responding will switch to the other stimulus for all subsequent trials. In learning set, this is the performance that is reinforced and this is the performance that the organism's behavior approaches. Thus, the correspondence between the reinforcement contingencies and the

behavior generated by these contingencies remains an appropriate criterion for this behavioral class.

Learning set does not simply demonstrate the complexity of the dimensions that the analysis of discriminated operants must take into account. The generalization of responding across problems or situations is a component of many important classes of behavior. For example, we call responding *imitative* when one organism duplicates the behavior modeled by another organism. Imitative responding is a class of responses that may be differentially reinforced (e.g., Baer, Peterson, and Sherman, 1967), and because it is occasioned by stimuli such behavior is an instance of a discriminated operant. We do not know the pertinent physical dimensions of imitative behavior (e.g., if touching the head is modeled in the game of "Simon says," the leader's hand is seen but not felt by the child whereas the child's own hand is felt but not seen; given that there is no simple correspondence between someone else's seen hand and one's own felt hand, how does the child learn to imitate?). In fact, we might ask what dimensions of correspondence control our response of calling certain classes of behavior imitative. When one organism learns something simply by watching another organism's behavior, we speak of *observational learning* (e.g., Myers, 1970), but the study of observational learning seems to differ from the study of imitation mainly in how explicitly previous training is arranged. In other words, we are most likely to call an instance of imitation observational learning when we observe such behavior but do not know how it originated. We should not even exclude the possibility that the capacity to learn by observing another organism does not itself depend on prior learning (e.g., Meltzoff and Moore, 1977).

The phenomenon of *learned helplessness* (S. F. Maier, Seligman, and Solomon, 1969) might also be regarded as a variation on learning set. For example, rats that receive inescapable and unavoidable electric shock in one situation will be less likely to learn avoidance responding when such responding becomes available in a new situation than rats not exposed to inescapable and unavoidable shock. Here again it is difficult to define the response class, except to note that the rats in the presence of inescapable and unavoidable shock apparently learned that their responses did not have important consequences, and the behavior generated by these contingencies transferred to situations in which responding could have consequences. The role of discriminative stimuli in this instance also is not obvious, though we might assume that through differential contingencies the behavior called learned helplessness could be brought under the control of discriminative stimuli.

Such examples illustrate the range and complexity of the stimulus and response classes that organisms can learn. Learning set also suggests that simple instances of learning cannot be observed simply by introducing an experimentally naive organism into an experimental setting. To study learning, therefore, some investigators have turned to the repeated acquisition of simple discriminations or simple response sequences. Paradoxically, they have come to study learning, defined by changes in performance, in the context of steady-state procedures, defined by stability in performance (Sidman and Rosenberger, 1967; Boren and Devine, 1968). For example, Boren and Devine (1968) arranged four groups of

three levers each within a monkey's chamber. A particular sequence of presses was reinforced within a given session (e.g., left lever of group 1, right lever of group 2, middle lever of group 3, right lever of group 4), but the sequence changed from session to session. After the monkey's rate of mastering a new problem each day becomes stable, the repeated acquisitions permit the study of a diversity of learning variables (e.g., the effects of different fading procedures, or of different types of consequences for responses at various positions in the sequence, or of establishing the response sequence as a whole as opposed to building it up either from the beginning or from the end).

SUMMARY

We began this chapter by exploring parallels between differentiation and induction on the one hand and discrimination and generalization on the other. Both are outcomes of differential reinforcement; they differ mainly in whether this operation is arranged for stimulus properties or response properties. Functional aspects of stimulus control were treated in the context of experiments on attention, on stimulus-control gradients, and on fading procedures. In dealing with discriminable properties of the environment, we considered how complex relational features of the environment could serve as discriminative stimuli. Examples included matching-to-sample, natural concepts, preparedness, and learning set. All of these cases emphasized stimuli, but we found it important to treat these stimuli in terms of their relations to responses and consequences. Stimulus control is based upon the three-term contingency: stimulus-response-consequence. None of these terms is significant in isolation. We may speak of differential reinforcement with respect to stimulus properties, or of the correlation of stimuli with contingencies, or of operant classes defined by the stimuli that set the occasion for responses in a class. These variations on the language of discrimination learning simply stress different aspects of the same three-term contingency.

Reinforcement Schedules

VARIABLE-INTERVAL AND VARIABLE-RATIO SCHEDULES

Yoked schedules

The role of interresponse times

Limited hold

Reinforcement schedules and causality

FIXED-INTERVAL AND FIXED-RATIO SCHEDULES

THE VOCABULARY OF REINFORCEMENT SCHEDULES

SCHEDULE COMBINATIONS

Multiple schedules: Behavioral contrast

Chained schedules: Conditioned reinforcement

Concurrent schedules: Matching, maximizing, and choice

SUMMARY: SCHEDULE ANALYSIS

EIGHT

Schedule is derived from the Middle English *sedule*, a slip of parchment or paper, which is in turn derived from the Latin *scheda*, papyrus leaf, and the Greek *skhizein*, to split. The Indo-European root, *skei-*, to cut or split, links *schedule* to *schizo-*, as in *schizophrenia*, and to *science* and *conscious*, from the Latin *scire*, to know, in the sense of being able to separate one thing from another.

Relatively few classes of responses have consistent consequences. The reinforcement of some responses but not others, sometimes called *intermittent* or *partial* reinforcement, is a general feature of behavior. For example, winning is not an invariable consequence of placing a bet, and finding a particular commodity is not an invariable consequence of going to a store, and getting an answer is not an invariable consequence of asking a question. Consider calling a friend on the telephone. Sometimes this response is reinforced by the opportunity to talk with one's friend. At other times, the friend does not answer or the friend's line is busy. *Continuous* or *regular* reinforcement, the reinforcement of every response within an operant class, is the exception rather than the rule. For this reason, we must examine the effects of *schedules of reinforcement*, arrangements that specify which responses within an operant class will be reinforced. We already have considered some schedules (ratio schedules in chapter 1, differential-reinforcement-of-low-rate schedules in chapter 6, and interval schedules in chapter 7). Schedules can arrange reinforcement on the basis of response number, the time at which responses occur, or the rate of responding. Number, time, and rate requirements can also be combined in diverse ways to produce more complex schedules. We will be concerned with the effects of reinforcement schedules not only as valuable

experimental tools, but also as ubiquitous and perhaps fundamental properties of behavior in their own right (cf. Ferster and Skinner, 1957; Schoenfeld and Cole, 1972).

Let us return to the telephone example. Suppose we call a friend and get no answer. The likelihood of getting an answer on later calls ordinarily depends on when we call and not on how many times we call. The friend will be available and thus the call will be answered only if it occurs at a particular time. Similarly, suppose we call and get a busy signal. Our calling again will not affect how long the busy signal lasts. Some variable time will elapse that depends only on how long the friend talks with someone else. Once the friend hangs up, our next call will produce a ring and an answer rather than a busy signal. We must call to get an answer, but again there will be an answer only if the call occurs at a particular time. These cases are everyday approximations to *variable-interval* or *VI* reinforcement schedules. The schedules arrange reinforcement for a single response that occurs after a specified time has elapsed, and this time varies from one instance to the next; responses before this time have no effect. A *VI* schedule is designated by the average time to the availability of a reinforcer. Because the times when people are away from a telephone are usually longer than the durations of telephone conversations, the situation in which a telephone rings but is not answered could be described as a longer-valued *VI* schedule than that in which the call produces a busy signal.

Now consider a different situation. We have to make a call at a pay telephone but do not have any change. We might ask passers-by for change, and therefore asking may be reinforced by getting the change we need for the call. In this instance, the likelihood that the response will be reinforced does not depend on the passage of time. Instead, we expect only a certain proportion of the people we ask to be able and willing to make change for us. We have to keep asking until we encounter such a person. We might be successful after asking only a very few people, or we might have to ask many. In other words, the availability of the reinforcer depends on number of responses rather than on the passage of time, and this number will vary from one reinforcer to the next. Such schedules are called *variable-ratio* or *VR* schedules. They are designated by the average number of responses required per reinforcer (i.e., the average ratio of responses to reinforcers).

These examples demonstrate that clearly specifying the response and the reinforcer is critical when we interpret behavior in terms of schedules. It is not sufficient to speak simply of telephoning. The behavior that leads up to calling, such as finding a pay telephone or getting change for it, is different from the calling itself. The circumstances of the call, such as whether an earlier call was made and whether that call resulted in no answer or a busy signal, must also be considered. Calling a friend who keeps on irregular schedule is different from calling a store or office that opens promptly for business at a particular time. Calling after a busy signal is different if the call is to someone who usually leaves the line free after calls than if it is to an active business at which lines do not remain open very long. And if the reinforcer is speaking to a particular person, we will analyze the situation

differently than if the reinforcer is some other consequence, such as a telephone sale made by someone in business or the recruitment of a volunteer for a fund drive. Getting an answer to a telephone call is usually reinforced according to a VI schedule, but because making a sale or recruiting a volunteer usually depends on the number of calls made, these events typically reinforce calling according to a VR schedule. We will not ordinarily have problems when we deal with simple responses such as pigeons' key pecks or rats' lever presses and simple reinforcers such as food or water deliveries. We will concentrate on such cases in the following survey. It is important to remember, however, that accurately applying the language of reinforcement schedules to settings outside the laboratory demands that we carefully specify the contingencies and the responses and reinforcers that enter into these contingencies.

VARIABLE-INTERVAL AND VARIABLE-RATIO SCHEDULES

In a variable-interval or VI schedule, the delivery of a reinforcer depends on the passage of a variable time and then the emission of a single response. In a variable-ratio or VR schedule, the delivery of a reinforcer depends on the emission of a variable number of responses without regard to the passage of time. Some properties of VI and VR schedules are illustrated by hypothetical cumulative records in Figure 8-1. Three records that might be generated by VI 1-min reinforcement are shown on the left (*A*, *B*, and *C*); three records that might be generated by VR 100 reinforcement are shown on the right (*D*, *E*, and *F*).

Consider first how we might arrange a VI schedule. We take a tape or film strip that can be driven at a constant speed by a motor and punch holes in it at irregular intervals. Using a switch that senses the location of these punched holes, we then allow a response to be reinforced each time one of these holes reaches the switch. A portion of such a tape is shown just above the time scale in Figure 8-1.¹ An alternative to the tape as a method for VI scheduling is to use a clock that provides a sequence of electrical pulses and randomly to select some proportion of these pulses to arrange or set up a reinforcer for the next response. Each method has advantages and disadvantages: With a tape, the experimenter must be concerned with the effects of repeating the sequence over and over; with random selection from a pulse source, the experimenter does not know ahead of time how long the next interval will be (cf. Fleshler and Hoffman, 1962; Farmer, 1963; Millenson, 1963; Catania and Reynolds, 1968, Appendix II).

Records *A*, *B*, and *C* of Figure 8-1 illustrate one of the most important

¹One technical feature of VI scheduling not incorporated into Figure 8-1 is whether the tape stops or continues running once a reinforcer has been scheduled for the next response. Although both arrangements are used, stopping the tape until the reinforcer has been delivered is the more common procedure, because it guarantees that intervals are timed from some environmental event, namely the reinforcer delivery, with which the organism has come in contact.




Figure 8-1. Hypothetical segments of cumulative records from performances maintained by a one-minute variable-interval schedule (VI 1-min) and a one-hundred-response variable-ratio schedule (VR 100). The horizontal strip above the VI time scale and the vertical strip next to the VR response scale show portions of punched tapes used to select responses for reinforcement. The three VI records, A, B, and C, show that these different response rates do not affect the rate of reinforcement but the number of responses per reinforcement does vary with response rate. The three VR records, D, E, and F, show that the rate of responding affects the rate of reinforcement but the number of responses per reinforcement is independent of response rate.

properties of a VI schedule: The rate at which reinforcers are delivered is not affected by the rate of responding. Even though the response rates in the three records are considerably different, responding produces the same number of reinforcers in each case. This schedule provides a specified and relatively constant rate of reinforcement over a substantial range of possible response rates. There is, of course, a lower limit to the response rate that will produce all of the scheduled reinforcers. Reinforcers are not delivered unless responses occur. Thus, reinforcers will be produced by responses less often than specified by the VI schedule if the time between successive responses becomes long relative to the time scheduled between one reinforcer and the next. With responses such as pigeons' key pecks reinforced by food, response rates are usually high enough in practice that this is not a serious problem. This independence of reinforcement rate from response rate, when coupled with the relatively constant response rate generated by the varying intervals from one reinforcer to the next, makes VI reinforcement a frequent choice of baseline schedule (e.g., for studying stimulus-control gradients: chapter 7).

Now consider arranging a variable-ratio or VR schedule. Again we use a punched tape or film strip, but this time we make the organism's responses rather than the passage of time move the tape relative to the switch. A portion of a tape for a VR 100 schedule is shown vertically next to the response scale accompanying

the VR records on the right in Figure 8-1. Each response advances the tape a fixed distance relative to the switch, and when the switch senses a hole in the tape the next response is reinforced. On the average, one response is reinforced per hundred responses but the number varies from one reinforcer to the next. In this case, number of responses per reinforcer is independent of response rate, and rate of reinforcement depends on response rate. In Figure 8-1, response rate is highest in *D*, and thus the reinforcers are delivered within a shorter period of time than in *E*; similarly, equal periods of time are represented in *E* and *F*, but more responses are reinforced in *E* than *F* because response rate is higher in *E* than in *F*.

As with VI schedules, alternative methods of scheduling are available. For example, a device that randomly selects some fraction of the emitted responses for reinforcement might be substituted for the tape. With either method of scheduling, the VR schedule like the VI schedule generates roughly constant rates of responding between reinforcers, at least with moderate ratio sizes. When the ratio of responses to reinforcers becomes very large, VR responding may be interrupted frequently by pauses.

Rates of responding maintained by VI and VR schedules are illustrated in Figures 8-2 and 8-3. Both figures show data obtained with pigeons' key pecks reinforced with food. In Figure 8-2, rate of pecking is plotted as a function of rate




Figure 8-2. Rate of key pecking as a function of rate of VI reinforcement, for six pigeons. The two scales at the bottom show the correspondence between rate of reinforcement and average interreinforcement intervals in VI schedules. (Adapted from Catania and Reynolds, 1968, Figure 1.)




Figure 8-3. Rate of key pecking as a function of VR schedule, for three pigeons. The two scales at the bottom show the correspondence between responses per reinforcement and the probability of reinforcement. (Adapted from Brandauer, 1958, Table 2.)

of reinforcement provided by the VI schedule (e.g., VI 1-min provides a maximum of 60 reinforcements per hour). Rate of pecking increases as a function of reinforcement rate, but the function is negatively accelerated, i.e., the increment in response rate produced by a given increment in reinforcement rate becomes smaller as reinforcement rate increases.

In Figure 8-3, rate of pecking is plotted as a function of the ratio, responses per reinforcement, arranged by the VR schedule. At VR 1 (leftmost data point), every response is reinforced. Even with reinforcement duration excluded, responding only slightly exceeded 25 responses per minute. Response rate was substantially higher at VR 10 (second data point), and reached its maximum, greater than 200 responses per minute, at VR 50. Response rate then decreased gradually with increasing VR size.

In their effects on behavior, VI and VR schedules differ considerably. For example, the different scales of the two figures show that over most of the range of values for each schedule VI response rates are lower than VR response rates. The schedules also differ in the way responding is affected when reinforcement is reduced or discontinued. The idealized cumulative records of Figure 8-4 show responding maintained by VI and VR reinforcement and responding during extinc-

tion after VI and after VR reinforcement. A low rate of VI reinforcement produces less responding than a high rate of VI reinforcement, but in both cases responding is fairly uniformly distributed in time; in addition, the decrease in response rate during extinction after VI reinforcement is gradual. With VR reinforcement, on the other hand, the decrease in response rate with larger ratios comes about in part because responding begins to be interrupted by long pauses (the appearance of long pauses during ratio performance is sometimes referred to as *ratio strain*); in addition, extinction after VR reinforcement usually produces abrupt transitions from high response rates to periods of no responding.

What properties of VI and VR schedules are responsible for their characteristically different performances? It seems reasonable that VR schedules should generate higher rates of responding than VI schedules, because an increase in VR response rate makes the delivery of reinforcers more frequent whereas an increase in VI response rate cannot have this effect. In what way can these different relations between response rate and reinforcement frequency act on an organism's behavior? Whether the schedule is VR or VI, the separation between successive




Figure 8-4. Temporal patterning of responding during VI and VR reinforcement and in extinction after VI or VR reinforcement. These somewhat idealized cumulative records contrast the effects of the two types of schedules: the relatively constant responding maintained by VI schedules and the gradual decreases in response rate during extinction after VI reinforcement, and the abrupt transitions between high rates of responding and long pauses in VR schedules with large response requirements and during extinction after VR reinforcement. Both types of schedule generate substantial quantities of responding in extinction.

reinforcers is variable. In either schedule, therefore, several closely spaced reinforcers will sometimes happen to be scheduled just as response rate decreases slightly, or a long period without reinforcement may happen to be scheduled just as response rate increases.

Yoked schedules

One method for studying the variables that operate within schedules is the yoked-chamber procedure (Ferster and Skinner, 1957), which arranges that an organism's performance in one chamber determines the events that occur in a second organism's chamber. The role of reinforcement rate in VR and VI schedules can be studied by arranging a VR schedule for one pigeon's key pecks so that each reinforcer produced by this pigeon schedules a VI reinforcer for the next peck of a second pigeon. In other words, the second pigeon's pecks are maintained by a




Figure 8-5. Cumulative records for the third session of yoking for two pairs of pigeons. The responses per reinforcement from Pigeon 402's VI performance were used to generate a yoked VR schedule for Pigeon 410's pecks. The interreinforcement times from Pigeon 414's VR performance were used to generate a yoked VI schedule for Pigeon 406's pecks. The horizontal arrows connecting the left pair of records show the correspondence of responses per reinforcement for that schedule pair; the vertical arrows connecting the right pair show the correspondence of interreinforcement intervals. In both cases of yoking, VR response rate was higher than VI response rate. A sample of the preliminary session that preceded sessions of yoking is shown in the inset for Pigeon 402 at the upper left. (From Catania, Matthews, Silverman and Yohalem, 1977, Figure 1)

VI schedule in which the successive interreinforcement intervals match those produced by the first pigeon's VR performance. In these circumstances, the two schedules differ in responses per reinforcer but not in time between successive reinforcers. Conversely, the role of responses per reinforcer can be studied by arranging one pigeon's VI schedule so that the responses emitted per reinforcer by this pigeon determine the ratios of a second pigeon's VR schedule. In this case, the second pigeon's pecks are maintained by a VR schedule in which the successive *ratios match those produced by the first pigeon's VI performance*. In these circumstances, the two schedules *differ in the times between successive reinforcers* but not in the responses per reinforcer.

Cumulative records from both types of yoking are illustrated in Figure 8-5 (Catania, Matthews, Silverman, and Yohalem, 1977). In one group of pigeons, pairs were assigned to yoked schedules in which one pigeon's pecks per reinforcement in a VI 30-second schedule were used to generate a second pigeon's VR schedule, as illustrated by Pigeons 402 and 410. In another group of pigeons, pairs were assigned to yoked schedules in which one pigeon's interreinforcement intervals in a VR schedule were used to generate a second pigeon's VI schedule, as illustrated by Pigeons 414 and 406. The representative pairs of Figure 8-5 show that, independent of whether VR is yoked to VI or VI is yoked to VR, VR schedules generate higher response rates than VI schedules (see also Zuriff, 1970). Furthermore, the difference between VR and VI rates emerges rapidly: the records are from the last of only three 50-reinforcement sessions.

The role of interresponse times

The yoking experiment shows that the rate difference between VR and VI schedules cannot be attributed solely to either responses per reinforcement or times per reinforcement. One possibility is that, even with yoking, reinforcement rate on the average changes with response rate in VR but not in VI schedules. It might be argued that the organism simply learns that in VR but not VI schedules faster responding produces more frequent reinforcement. An alternative is the relation between interresponse times (IRTs) and the probability of reinforcement (e.g., Anger, 1956; Shimp, 1973; Alleman and Platt, 1973; Kuch and Platt, 1976). An interresponse time or IRT is simply the time between two successive responses (cf. DRL schedules in chapter 6). In VR schedules, the probability that a response will be reinforced depends only on the ratio; it does not vary with the time elapsed since the preceding response. In VI schedules, the probability that an interval has ended elapsed since the last response will be reinforced. In other words, a larger proportion of long than short IRTs is reinforced in VI than in VR schedules, and therefore relative to VR schedules VI schedules differentially reinforce long IRTs. An increase in these longer times between responses necessarily implies lower response rates.

These relations are illustrated in Figure 8-6. Some event records from hypothetical VI and VR performances are shown at the left; vertical lines represent




Figure 8-6. Hypothetical event records from VI and VR performances, on the left, illustrate the derivation of functions relating probability of reinforcement to IRT, on the right. In the event records, vertical lines represent responses, arrows represent reinforcers, and IRTs are numbered consecutively.

responses and arrows represent reinforcers. In the three VI records (top), rate of responding does not affect the time at which a response becomes eligible for reinforcement. In each case a reinforcer is delivered for the response at the end of the 5-second segment. When responses occur once per half-second (first record), the last of ten half-second IRTs is reinforced, i.e., half-second IRTs are reinforced with a probability of 0.1. When responses occur once per second (second record), the last of five one-second IRTs is reinforced, i.e., one-second IRTs are reinforced with a probability of 0.2. Finally, when responses occur once per 2.5 seconds (third record), the last of two 2.5-second IRTs is reinforced, i.e., 2.5-second IRTs are reinforced with a probability of 0.5. In the three VR records (bottom), however, time between responses does not affect probability of reinforcement. In each case, the last of five IRTs is reinforced. The probability of reinforcement for an IRT is therefore 0.2 whether responses occur at four or two or one per second. Functions relating probability of reinforcement to IRT are shown for VI and VR schedules on the right in Figure 8-6.

The higher probability of reinforcement after long than short IRTs implies that VI schedules maintain lower responses rates than VR schedules because VI schedules differentially reinforce long IRTs. Yet the relatively slow development of temporally spaced responding when such differential reinforcement is explicitly scheduled (e.g., DRL schedules, chapter 6) raises questions about such an account (see also Dews, 1962; H. M. Jenkins, 1970; Catania, 1971). An alternative is that organisms are sensitive to the distributions of reinforced and unreinforced IRTs. The relation between these distributions is different with VR than with VI schedules. Because the probability of reinforcement is independent of IRT in VR

schedules, the distribution of reinforced IRTs has the same form as the distribution of all IRTs. This is not so with VI schedules; because the probability of reinforcement increases with IRT, a relatively larger proportion of long IRTs is reinforced and therefore the distribution of reinforced IRTs is shifted toward longer IRTs relative to the distribution of all IRTs (cf. Revusky, 1962). These relations among IRT distributions are inevitable features of VR and VI scheduling. Organisms must therefore in some sense be sensitive to them when their responses are affected by the respective schedules.

Limited hold. One temporal contingency sometimes added to schedules is called the *limited hold* or *LH*. With a limited hold, a scheduled reinforcer remains available or is held for a response only for a limited time; if no response occurs within that time, the reinforcer is lost. The telephone again is an illustration. If one gets a busy signal when calling into an overloaded switchboard, getting a ring and an answer may not become more likely as time passes, because the lines never remain open very long. If the lines at such a switchboard are used again within only five seconds after someone hangs up, we describe the schedule as a VI schedule with a five-second limited hold. A caller is most likely to get through such a switchboard by hanging up immediately after a busy signal and dialing again within five seconds. In the laboratory, adding a limited hold to a schedule typically increases response rate (e.g., Hearst, 1958); in general, the shorter the limited hold the greater this increase, except that a very short limited hold may allow so few responses to be reinforced that the schedule fails to maintain responding. The effect of the limited hold on VI performance can be interpreted in terms of its attenuation of the differential reinforcement of long IRTs: the probability of reinforcement does not increase as IRT durations approach the duration of the limited hold, because the limited hold may cancel the most recently available reinforcer before the next response occurs.

Reinforcement schedules and causality

The effects of reinforcers depend on the responses they follow, but reinforcers can follow responses both when produced by responses and when delivered independently of responses. Is responding affected in the same way when it produces a reinforcer as when it happens by accident to be followed by a reinforcer? We gave this question a preliminary answer when we considered the phenomenon called superstition. The role of the temporal contiguity between responses and reinforcers has been discussed by Herrnstein (1966):

Is it possible that the accidental correlations in time among responses, stimuli, and reinforcers do not exert control over behavior? . . . One of the characteristics of accidental correlations between behavior and environmental events is *variability*. Every aspect of behavior may vary and yet be contiguous with a reinforcer that is independent of the behavior. In contrast, behavior that is instrumental must have at least one aspect that has a more or less fixed

correlation with the reinforcer. Were animals sensitive to this difference, they could detect those events over which their behavior has no real control. (Herrnstein, 1966, pp. 42-43)

Consider now an experiment by Lattal (1974). Pigeons' key pecks were reinforced according to a VI schedule. Once VI performance was well established, the schedule was changed: at the end of some intervals, the reinforcer was delivered immediately, without regard to responding. For example, when only 33% of the reinforcers were response-produced, a peck produced food at the end of every third interval but food was delivered independently of pecks in the remaining intervals. With decreasing percentages of response-produced reinforcers, the rate of responding decreased. When the percentage of response-produced reinforcers was zero, so that food was completely independent of behavior, rates of responding approached zero (this arrangement is equivalent to a superstition procedure with variable times between reinforcer deliveries; it is sometimes called a variable-time or VT schedule; Zeiler, 1968). The decrease in response rate has interesting implications. For example, when 33% of the reinforcers were response-produced, response rates were roughly half those when all reinforcers were response-produced, but even the pigeon with the lowest response rate responded at about thirty responses per minute, or one response every two seconds. This implies that response-independent reinforcers followed responses on the average by only about one second (because the average location of the reinforcer would be in the middle of a two-second IRT). Given, therefore, that every reinforcer was likely to follow closely after some response, and that one-third (the response-produced reinforcers) followed responses immediately, why did the accidental temporal relations between responses and reinforcers not maintain responding at a rate close to that when all reinforcers were response-produced?




Figure 8-7. Hypothetical segments of event records from VI, VI with delay, and VT reinforcement schedules (left), and three pigeons' rates of key pecking maintained by these schedules (right). (Adapted from Sizemore and Lattal, 1977, Table 1.)

When reinforcers are delivered independently of responses, the time between the most recent response and the reinforcer is likely to vary from one reinforcer to the next. Thus, Lattal's data suggest that this variability does counteract superstitious responding. But what about cases in which the time between the most recent response and the reinforcer varies even though the reinforcer is response-produced? This occurs, for example, when delay of reinforcement is added to a schedule. A case is illustrated in the event records in Figure 8-7 (left). Vertical lines represent responses and arrows represent reinforcers. The top record shows a segment from a standard VI schedule: the interval ends at the dashed line, and the next response, *a*, is followed immediately by the reinforcer. A segment from a VI schedule to which a three-second delay of reinforcement has been added is shown in the middle record: the interval ends at the dashed line, and the response *b* produces a reinforcer three seconds later; because other responses occur during this time, the time between the last response and the reinforcer (*c*) is shorter than the delay interval (*d*), and this time varies depending on the spacing of responses during the delay.² The bottom record shows a segment from a VT or response-independent schedule: when the interval ends, at the dashed line, the reinforcer is delivered, and because it occurs independently of responses, the time between the last response and the reinforcer (*e*) varies.

The bar graph in Figure 8-7 shows data from three pigeons obtained with these schedules (Sizemore and Lattal, 1977; see also B. A. Williams, 1976). The rates of pecking maintained by VI with delay were between the moderate rates maintained by the VI schedule and the low rates maintained by the response-independent reinforcers of the VT schedule. We just concluded that the difference between VI and VT must depend in some way on the variable times between the last response and the reinforcer (*c* and *e*). Why then are the rates with delayed VI reinforcement higher than those with VT reinforcement? The issues remain open to experimental analysis, but it is reasonable to assume that when response rate decreases with delayed VI reinforcement, it eventually reaches a point at which pauses longer than the delay occur often enough that sometimes no other responses intervene between the reinforcer and the response that produced it. In these instances the times between the last response and a reinforcer can never be longer than the scheduled delay, and therefore they will no longer be variable. With VT reinforcement, on the other hand, there is no upper limit on these times, and therefore they must always be variable.

Experiments like these show how sensitive organisms are to the consequences of their own behavior. This sensitivity seems to depend not merely on the closeness of responses and consequences in time but also on the distributions of the times that separate responses and reinforcers. These distributions distinguish events

²Delay of reinforcement was discussed in chapter 5 (cf. Figure 5-5). Other procedures for studying reinforcement delay include the presentation of a distinctive stimulus during the delay, which is an immediate consequence of responding that may acquire reinforcing properties of its own through its relation to the later reinforcer delivery, and the requirement that no responses occur during the delay, when the procedure is equivalent to the differential reinforcement of a response followed by a pause (cf. Ferster, 1953; Dews, 1960).

caused by behavior from events accidentally correlated with behavior; they are therefore central to the concept of causality. For this reason, the study of reinforcement schedules is fundamental, and the demonstration that humans are similarly sensitive to the consequences of their own behavior has both theoretical and practical significance (cf. Holland, 1958; Weiner, 1969; B. A. Matthews, Shimoff, Catania, and Sagvolden, 1977; Rachlin and Frankel, 1969).

FIXED-INTERVAL AND FIXED-RATIO SCHEDULES

If the probability that a response will be reinforced is greater at some times than at others, we can expect the rate of responding to be higher at those times than at the others (Catania and Reynolds, 1968). For example, suppose most intervals in a VI schedule are ten seconds long, and the remaining intervals are between fifty and one hundred seconds. Response rate will probably be high at about ten seconds after reinforcement, and if a response is not reinforced at ten seconds response rate might decrease over the next twenty or thirty seconds before increasing again later on. We have mostly considered VI and VR schedules designed to hold the probability of reinforcement roughly constant over time (in VI) or over number of responses (in VR). But schedules can be arranged in which time to the availability of a reinforcer or number of responses per reinforcer is constant from one reinforcer to the next; such schedules are respectively called *fixed-interval* or *FI* schedules and *fixed-ratio* or *FR* schedules (in addition, schedules in which the time between response-independent reinforcer deliveries is constant are called *fixed-time* or *FT* schedules). The fixed schedules are of interest because they introduce discriminable periods during which no reinforcers occur.

Consider first the fixed-interval or FI schedule: a response is reinforced only after some constant time has elapsed since some environmental event; responses before this time have no effect. An example is looking at one's watch as the end of a lecture period approaches; in this case, the reinforcer is seeing that the time has come at which one can inconspicuously leave the classroom. Looking at the watch more often has no effect: it does not make the watch run more rapidly. The characteristic property of the FI schedule is that responding tends to occur at zero or low rates early in the interval and to increase as the end of the interval approaches. We might similarly expect that the frequency with which one looked at one's watch would increase as the end of a lecture period approached.

Sample cumulative records of FI performance are shown in *B* and *C* of Figure 4-6 (chapter 4). Each record shows a sequence of fixed intervals. The concave-upward pattern of such cumulative records is sometimes called *FI scalloping*. Depending on the sessions of exposure to FI schedules and other variables, individual FI scallops may show relatively abrupt transitions from no responding to a roughly constant rate of responding, as in *B*, or gradually increasing rates after responding begins, as in *C*. The pattern of FI responding tends to be consistent over relative rather than over absolute time in the interval. For example, if respond-

ing reaches half its final or terminal rate 40 seconds into a 100-second fixed interval, it is likely to do so 20 seconds rather than 40 seconds into a 50-second fixed interval. (This consistency in relative rather than absolute time appears to be a property of elicited as well as consequential behavior; cf. Killeen, 1975.)

Any theoretical treatment of FI performance must consider the finding that the FI scallop survives repeated interruptions. The phenomenon is illustrated in Figure 8-8 (Dews, 1962). The key pecks of four pigeons were reinforced with food according to FI 500-second schedules. The left graph shows the average rate of pecking in successive fifty-second portions of the interval. In another procedure, the houselight, which provided general illumination in the chamber, was turned on and off in successive fifty-second portions of the interval; after the last fifty seconds of the interval, it remained on until the peck at the end of the interval was reinforced. Pecking was not reinforced in the absence of the light, and pecking during these times decreased in rate, as shown by the dark bars in the right graph of Figure 8-8. Nevertheless, the increase in responding in successive periods when the houselight was on, as shown by the shaded bars, was similar in form to the increase when there were no interruptions of FI responding during the interval. Because of findings such as these, the FI performance has received varied theoretical treatments (e.g., Gollub, 1964; Zeiler, 1977; B. A. Schneider, 1969; Dews, 1970b; Hawkes and Shimp, 1975). Among the issues are whether the FI scallop represents a gradient of delayed reinforcement (in that responses at various locations in the interval are consistently followed by the reinforcer at the end of the interval) or a gradient of temporal discrimination (in that the organism's different rates at different times in the interval imply discrimination of the elapsed time) or some combination of such gradients.

Let us now turn briefly to fixed-ratio or FR schedules. In these schedules, the last of some fixed number of responses is reinforced. For example, it usually




Figure 8-8. Rate of pecking in successive fifty-second periods during a standard FI 500-second schedule of reinforcement (open bars, left), and during an FI 500-second schedule in which a light (shaded bars) and its absence (dark bars) alternated every fifty seconds (right), for four pigeons. The periods when the light was absent produced low response rates and therefore interrupted FI responding, but the rate when the light was present increased in much the same way as in a standard FI. (Adapted from Dews, 1962, Figure 2.)

requires a fairly constant number of turns of the winding key to wind up a spring-driven device such as a toy car or a music box. Reaching the point at which the key can no longer be turned is ordinarily a reinforcing consequence, because it sets the occasion for releasing the car and watching it run or for listening to the tune played by the music box. An FR schedule arranges that after a reinforcer has been delivered at the end of one ratio, no response will be reinforced until another ratio has been completed. Because the first response after a reinforcer delivery is never reinforced, FR schedules typically generate a performance characterized by postreinforcement pauses followed by high-rate responding. The average duration of the postreinforcement pause increases with increasing FR size (e.g., Felton and Lyon, 1966).

The FR performance is sometimes called break-and-run responding; once responding begins after the postreinforcement pause, it typically occurs at a high rate and without interruption until the reinforcer is delivered. The cohesiveness with which FR responding is maintained once it begins has suggested that the FR sequence might be regarded not simply as a succession of responses but as a behavioral unit in its own right. Ordinarily, the development of FR responding proceeds so rapidly that its details cannot easily be seen. Figure 8-9, however, shows a gradual development of FR responding. In a two-key pigeon chamber, the




Figure 8-9. Development of FR 100 performance for Pigeon 91's food-reinforced key pecks. The typically rapid development of FR responding was slowed down by the concurrent operation of a VI schedule of reinforcement (not shown) for pecks on a second key. The cumulative record segments are from early portions of the numbered sessions of FR reinforcement.

pigeon's pecks on one key were reinforced according to an FR 100 schedule while a VR schedule operated concurrently on the other key. The VI schedule operating on the other key retarded the development of the typical FR break-and-run performance, and in effect made it possible to examine this development in slow motion; the change in the temporal patterning of responses that usually takes place in two or three sessions took place over 150 sessions. In the early sessions, responding was fairly uniformly spaced between reinforcers. Lower response rates and eventually pauses began to follow reinforcers only with continued exposure. The portion of the ratio consisting of uninterrupted high-rate responding gradually became longer, so that in later sessions responding continued with few if any interruptions after each postreinforcement pause. With successive sessions, the FR run appeared to build itself up backwards from the reinforcer.

The treatment of schedule performances as response units in their own right led to the development of higher-order schedules of reinforcement. For example, if reinforcement depended on the completion of a certain number of fixed intervals, we could speak of this arrangement as a second-order schedule in which successive FIs were reinforced according to an FR schedule (e.g., Findley, 1962; Kelleher and Gollub, 1962; Bigelow, 1971). An analysis of the component performances in such schedules would be concerned with the structure of behavior, and would be analogous to examining the properties that define an operant class (cf. chapter 6 and Shimp, 1976).

THE VOCABULARY OF REINFORCEMENT SCHEDULES

We have considered a variety of reinforcement schedules. The vocabulary for these schedules provides major classifications, but also includes admittedly idiosyncratic usages. For example, there is no particular rationale for distinguishing between FI and VI versus FT and VT schedules except that two different names were needed (we could as easily justify such names as fixed-duration and variable-duration, presumably abbreviated FD and VD). Nevertheless, it is important to note the distinctions. For example, FI, DRL and FT schedules all require the passage of a constant time, but they differ in response requirements. An FI schedule imposes no restriction on responding during the interval, but one response must occur at the end of the interval; some time must elapse, during which few or many responses can occur, and at the end of that time the next response is reinforced. A DRL schedule, however, requires a specified time without responses, and then the next response is reinforced; responses that occur too soon start the time over. Finally, in an FT schedule, the reinforcer is delivered at the end of the specified time whether or not responses have occurred. These schedule names were decided somewhat arbitrarily in the evolution of research on reinforcement schedules, but are well established (cf. Catania, 1968).

Although there have been several noteworthy attempts to classify reinforcement schedules more systematically (e.g., Schoenfeld, Cumming, and Hearst, 1956;

TABLE 8-1

Summary of Basic Reinforcement Schedules (T sec = time in seconds; N = number of responses)

Name and Abbreviation		Contingency	Comment
Variable interval	VI	T sec, then 1 response	T varies; with properly distributed intervals, generates roughly constant response rate.
Fixed interval	FI	T sec, then 1 response	T is constant; generates FI scallops
Variable ratio	VR	N responses	N varies; high constant rates, but large N produces breaks in responding called ratio strain.
Fixed ratio	FR	N responses	N is constant; generates postreinforcement pauses and high-rate runs.
Variable time	VT	T sec	T varies; provides response-independent reinforcers.
Fixed time	FT	T sec	T is constant; provides response-independent reinforcers.
Continuous reinforcement	(FR 1)	1 response	All responses reinforced; formerly abbreviated as CRF.
Extinction	EXT	—	As description of procedure, often used without regard to whether response had prior history of reinforcement.
Limited hold	LH	Reinforcer canceled if no reinforced response within T sec	T is constant if not otherwise specified; LH cannot stand alone, but must always be a feature added to some schedule.
Differential reinforcement of low rate (or long IRT)	DRL	T sec without response, then 1 response	Maintains responding easily because decreases in response rate increase reinforcement and thus prevent extinction.
Differential reinforcement of high rate	DRH	1 response within T sec or less of last response	Alternatively, at least N responses within T sec; responding may be difficult to maintain, because decreases in response rate lead to decreases in reinforcement.

TABLE 8-1 (continued)

Name and Abbreviation		Contingency	Comment
Differential reinforcement of paced responding	DRP	1 response between T and T' sec of last response	Sets both upper and lower limits on reinforceable response rates.
Differential reinforcement of other behavior	DRO	T sec without response	A negative-punishment or omission procedure; ordinarily decreases the likelihood of the designated response.

Skinner, 1958; Mechner, 1959; Dews, 1960; Snapper, Knapp, and Kushner, 1970; Schoenfeld and Cole, 1972), none has attained general usage. Table 8-1 summarizes the major procedural distinctions among schedules. The schedule definitions apply whether reinforcers are arranged successively and without interruption or occur within separate trials (e.g., an FI is ordinarily timed from the last reinforcer, but if other events were arranged between successive intervals the timing could begin with the onset of some stimulus, such as a color presented on a pigeon key).

The first two columns of Table 8-1 provide schedule names and their standard abbreviations. In practice, the abbreviations are usually accompanied by designations of time or number (e.g., VI 30-sec, FT 60-sec, LH 5-sec, DRL 10-sec, and FR 50). The next column describes the schedule contingency, or the conditions under which the schedule reinforces responses or delivers reinforcers (cf. especially the relations between FI and FT and between DRL and DRO). The last column provides brief commentary on each schedule. The vocabulary of Table 8-1 is in terms of reinforcement schedules, but can also be extended to punishment schedules (e.g., Azrin, 1956; Hendry and Hendry, 1963). The symmetry of reinforcement and punishment effects, illustrated in Figure 5-1, applies to scheduling effects as well (e.g., superimposing an FI schedule of punishment on maintained responding produces an inverted scallop, a gradually decreasing response rate as the end of the interval approaches).

SCHEDULE COMBINATIONS

Schedules need not operate in isolation. New schedules can be created without limit by combining the basic schedules in various ways (e.g., Herrnstein and Morse, 1958; Berryman and Nevin, 1962; Powers, 1968; Barrett, 1975, 1976). For example, reinforcers can be arranged for completing *either* of two schedule requirements (*alternative* schedules) or *both* of two schedule requirements (*conjunctive* schedules). In an alternative FI 30-sec FR 50 schedule, either the first response after thirty seconds or the fiftieth response will be reinforced, whichever comes first; in a conjunctive FI 100-sec FR 20 schedule, a response will not be reinforced until both 100 seconds have passed and nineteen responses have already been emitted. When time and number requirements interact, the arrangement is called an *interlocking* schedule. For example, an interlocking FR FI schedule might decrease

the duration of the interval as a function of number of responses, or perhaps increase the size of the ratio as a function of the passage of time (an instance of an interlocking schedule is winding a wristwatch, in which the reinforcer is the tension of the fully wound spring; to the point at which the watchspring is completely run down, the number of turns required to wind the watch increases with the passage of time).

The effects of schedule combinations are sometimes paradoxical. For example, Herrnstein and Morse (1958) found that an FI 15-minute schedule of food reinforcement for a pigeon's pecks maintained more than 300 pecks per interval, on the average. A response requirement was then added by changing to a conjunctive FR FI schedule. Even at a ratio requirement of FR 40, well below the number of responses that the pigeon emitted per interval anyway, the rate of pecking decreased. At larger ratios, up to a conjunctive FR 240 FI 15-minute schedule, response rate became so low that a pigeon occasionally took hours to produce a reinforcer that could have been available after only 15 minutes. This finding, that responding emitted when not a requirement of the reinforcement contingencies can sometimes be reduced by making it a requirement (see also Zeiler, 1977), demonstrates the challenge and variety of the analysis of reinforcement schedules.

Schedules may operate in sequence, alternating with each other; they may operate at the same time, either for the same response or for different responses; the consequence of completing one schedule may be the onset of another schedule; and different schedules may be accompanied either by the same stimulus or by different stimuli. These variables define several classes of schedule combinations, outlined in Table 8-2. Some schedules bear on various problems of historical significance, such as discrimination learning (multiple schedules), conditioned reinforcement (chained schedules), and choice (concurrent schedules). We will provide only a brief overview of these three areas in a sampling of specific research issues.

Multiple schedules:

Behavioral contrast

We considered multiple schedules under the topic of stimulus control (e.g., Figure 7-1 illustrates a multiple VI EXT schedule). We demonstrate stimulus control when two different schedules alternate, each in the presence of a different stimulus, and the performance appropriate to each schedule occurs during the corresponding stimulus. For example, if an FI schedule operates for a pigeon's pecks in green and a VI schedule for pecks in red, FI scallops in green may alternate with roughly constant response rates in red (note therefore that stimulus control can be demonstrated not only with reinforcement versus extinction but also with two different reinforcement schedules).

The performances in two components of a multiple schedule are not necessarily independent of each other. For example, if VI reinforcement maintains

TABLE 8-2

Compound Schedules. For convenience, each case is defined in terms of only two arbitrary component schedules, *A* and *B*, but compound schedules can include any number of components. In some instances, stimuli are designated by *S* with a superscript that indicates the schedule each stimulus accompanies.

<i>Schedule</i>	<i>Definition</i>	<i>Example (with Abbreviation)</i>
Alternative	Reinforcer is contingent on either <i>A</i> or <i>B</i> .	Completion of either a VR or a VI is reinforced, whichever occurs first (altern VR VI)
Conjunctive	Reinforcer is contingent on both <i>A</i> and <i>B</i> .	Reinforcer depends on the completion of both an FR and an FI (conjunc FR FI)
Interlocking	Reinforcer is contingent on some combined function of <i>A</i> and <i>B</i> .	Reinforcer is delivered when the sum of responses plus elapsed seconds equals some constant (interl FR FT)
Multiple	<i>A</i> during <i>S</i> ^A alternates with <i>B</i> during <i>S</i> ^B .	VI during red alternates with EXT during green (mult VI EXT)
Mixed	<i>A</i> alternates with <i>B</i> .	DRL alternates with FI, with no correlated stimuli (mix DRL FI)
Chained	During <i>S</i> ^A , completion of <i>A</i> produces <i>S</i> ^B ; during <i>S</i> ^B , completion of <i>B</i> produces reinforcer.	Completion of VR during blue produces yellow; completion of FR during yellow produces food (chain VR FI)
Tandem	Completion of <i>A</i> produces <i>B</i> ; completion of <i>B</i> produces reinforcer.	Completion of VR produces DRH, and completion of DRH produces food, with no correlated stimuli (tand VR DRH)
Concurrent	<i>A</i> operates for one response and <i>B</i> operates for another response.	One VI operates for pecks on left key and another VI for pecks on right key (conc VI VI)
Conjoint	<i>A</i> and <i>B</i> operate at the same time but independently for a single response.	VI and DRP operate simultaneously for presses on a single lever (conjnt VI DRP)
Second-order	Completion of <i>A</i> is reinforced according to <i>B</i> . (Reinforcing the second-order according to a third schedule, <i>C</i> , creates a third-order schedule, and so on for other higher orders.)	Successive FRs are treated as response units reinforced according to an FI schedule (FI [FR])

a pigeon's key pecking during one stimulus and the schedule in the presence of a second stimulus changes from VI reinforcement to EXT (extinction), decreased key pecking during the second stimulus is often accompanied by increased key pecking during the first. The phenomenon, called behavioral contrast, is illustrated in Figure 8-10 (G. S. Reynolds, 1961c). A multiple VI 3-minute VI 3-minute schedule operated for a pigeon's key pecks in red and green. The VI reinforcement in red was maintained while the schedule in green was changed from VI to EXT and then back to VI. The response rate in red increased during EXT in green, even though there was no change in the red VI schedule; both rates returned toward their earlier values when the green VI schedule was reinstated.

Behavioral contrast has received extensive analysis, ranging from accounts in terms of quantitative properties of schedules (e.g., G. S. Reynolds, 1961a; Lander and Irwin, 1968; Herrnstein, 1970; Rachlin, 1973) to accounts that appeal to frustrative or emotional effects of reduced reinforcement (e.g., Amsel, 1958; Terrace, 1968). Varieties of contrast effects have been distinguished, ranging from relatively permanent or sustained increases to those that last for only seconds or minutes after a schedule change (e.g., Catania and Gill, 1964; Nevin and Shettleworth, 1966; Staddon, 1969). Contrast effects also vary with responses, reinforcers, and organisms (e.g., Hemmes, 1973; Beninger and Kendall, 1975).

Empirical questions have been raised by behavioral contrast (e.g., whether it




Figure 8-10. Effects of extinction in one component of a multiple schedule on responding maintained by reinforcement in another component. The multiple-schedule components consisted of alternate three-minute periods of red and green projected on the pigeon key. In the sessions shown, a VI 3-min schedule of food reinforcement for the pigeon's pecks operated during red; the schedule during green was changed from VI 3-min to extinction and back to VI 3-min. The increased response rate in the unchanged component during the decreased reinforcement in the other component is called behavioral contrast. (After G. S. Reynolds, 1961c, as adapted by Terrace, 1966, Figure 10.)




Figure 8-11. Two interpretations of behavioral contrast: an excitatory effect on response 1 produced by extinction of response 2 (A), or an inhibitory effect on response 1 produced by reinforcement of response 2 (B). (From Catania, 1969a, Figure 4.)

is related to peak shift: see chapter 7), but the phenomenon also involves conceptual issues. Assume response 1 is a key peck in red and response 2 is a key peck in green. Both A and B in Figure 8-11 might be discussed in terms of behavioral contrast: response 1 (solid line) increases from its baseline level (X) to a higher level (Y) during extinction of response 2 (dashed line), and returns to its baseline level when reinforcement of response 2 is reinstated. But how are we to tell which is effect and which is baseline? We could assume equally well either that extinction of response 2 enhanced response 1, as shown by the arrows in A, or that

reinforcement of response 2 inhibited response 1, reducing it from its baseline level when reinforced by itself, as shown by the arrows in *B*.

The data argue that behavioral contrast is consistent with *A* of Figure 8-11, and can best be interpreted as the summation of two kinds of pecks: those maintained by reinforcement and those engendered by any stimulus that is differentially correlated with reinforcer deliveries (e.g., as when reinforcers occur more often during the VI than the EXT component of a multiple VI EXT schedule). The former class of pecks is an operant class, but the latter class is perhaps best treated as elicited responding. The two classes have different durations and topographies, and therefore can be separated experimentally (K. Keller, 1974; Schwartz, Hamilton, and Silberberg, 1975). For example, pigeons peck at the VI stimulus of a multiple VI EXT schedule even if the stimulus is on a second key at which pecks never produce food; this pecking occurs in addition to the pecks maintained on the first key by food reinforcement, and does not occur during a multiple schedule with equal VI components (K. Keller, 1974). Such pecks will be treated further in the topic of autoshaping in chapter 9.

Chained schedules:

Conditioned reinforcement

Chained schedules, introduced in chapter 1 to illustrate some apparatus and techniques in the analysis of behavior, have been used extensively to study *conditioned reinforcers*, reinforcers that acquire their capacity to reinforce through their relation to other stimuli already effective as reinforcers (e.g., Kelleher and Gollub, 1962; Zimmerman, Hanford, and Brown, 1967). For example, the feeder light becomes a reinforcer only through its relation to food in the feeder. Because the opportunity to engage in a highly probable response can reinforce a less probable response only if the opportunity is signalled in some way, the concept of conditioned reinforcement is intimately related both to the relativity of reinforcement and to the discriminative functions of stimuli (cf. Egger and Miller, 1962; Bower, McLean, and Meacham, 1966; Hendry, 1969; Catania, 1975; Fantino, 1977).

The example in chapter 1 used chained ratio schedules; Figure 8-12 uses chained interval schedules. The top record shows a chain of six successive FI 15-sec schedules arranged for Pigeon 64's key pecks, each in the presence of a different key color (the six components are numbered in the first chain of the record). In each component, the first response after fifteen seconds changed the key color to that of the next component of the chain (yellow to violet to white to green to red to blue), until the end of the last component when the final response produced the food reinforcer. The record, obtained after only three hours of the chained schedules, shows extensive pauses in the early components; the pauses became even longer with continued training. The middle record shows a segment of Pigeon 178's performance with corresponding tandem schedules. The only difference was that the tandem schedules operated in the presence of a single stimulus, vertical lines projected on the key. Pigeon 178 completed successive tandem sequences in little




Figure 8-12. Cumulative records of the performances of three pigeons on chained and tandem schedules. Both consisted of six consecutive FI 15-sec schedules in the last of which a peck produced food. The six components, numbered in the top record, are separated by pen displacements; food deliveries were followed by a record, are separated by pen displacements; food deliveries were followed by a reset of the pen to its baseline position. (Catania, Yohalem, and Silverman, unpublished data.)

more than the 90-second minimum demanded by the six successive 15-second intervals (the schedule differs from FI 90-sec in that at least six pecks are required, the schedule differs from FI 90-sec in that at least six pecks are required, intervals (the schedule differs from FI 90-sec in that at least six pecks are required, one at the end of each of the six 15-sec intervals). The bottom record shows the same chained and tandem schedules operating in alternation for Pigeon 392's pecks. Again long pauses occurred in the early components of the chain whereas the tandem sequence was completed expeditiously.

The chained-schedule stimuli reduced responding relative to the tandem schedule. How can we reconcile this phenomenon with the assumption that the successive stimuli of the chain should become conditioned reinforcers through their relation to the food at the end of the sequence? The presence of different stimuli in the chained schedule made a performance otherwise executed in a little more than 90 seconds sometimes take 15 minutes or more (with a corresponding decrease in the frequency of reinforcement). The low rates in the early components of chained schedules represent the combined discriminative effects of the chain stimuli, in that responding is never reinforced with food in the presence of these stimuli, and the reinforcing effects of the contingency between responding and the onset of the next stimulus. The data show that only the last components of the chain, close to the reinforcer, acquire substantial reinforcing effects.

The phenomenon is highly reliable (Kelleher, 1966; Gollub, 1977), and depends on a constant ordering of the chained stimuli; the long pauses are markedly reduced if the stimulus order changes from one reinforcer to the next (Kelleher and Fry, 1962). Analogous but inverse effects also hold for schedules of punishment; i.e., relative to tandem schedules, chained schedules of punishment mostly reduce responding only in the final components of the chain (Silverman, 1971; cf. Sidman and Boren, 1957). An implication is that punishment after a deed is done will probably have its greatest effect on the behavior that precedes getting caught and only minimal effects on the behavior leading up to the misdeed.

In chained schedules, we considered arrangements that attenuate the effects of consequences. Other schedules can be arranged in which brief presentations of stimuli greatly amplify the responding maintained by reinforcers (e.g., Findley, 1962; Stubbs and Cohen, 1972). For example, Findley and Brady (1965) reinforced a chimpanzee's push-button presses with food pellets according to an FR 4000 schedule. With no other consequences arranged, postreinforcement pauses ranged from many minutes to hours. But when the hopper light that accompanied the delivery of the pellets was briefly operated after every 400 responses of the FR schedule, responding increased substantially and postreinforcement pauses were typically five minutes or less. Stimuli that accompany reinforcers can be used to augment the effectiveness of those reinforcers. This example illustrates again that the effects of schedules depend critically on the detailed relations among stimuli and responses and consequences (cf. Morse and Kelleher, 1977).

Concurrent schedules:

Matching, maximizing, and choice

The behavior of organisms is not restricted to single responses. We must therefore examine the effect on one response of a reinforcement schedule operating for another response (cf. Catania, 1966). Concurrent schedules are schedules arranged simultaneously for two or more responses. Consider an FR 25 schedule of food reinforcement for a pigeon's pecks on one key, and an FR 50 schedule operating concurrently for pecks on a second key. By itself, either schedule maintains responding. When they operate concurrently, however, responding is likely to be maintained exclusively on the key with the FR 25 schedule (e.g., Herrnstein, 1958). The outcome should not be surprising. Each reinforcer requires only 25 pecks on the first key but it requires 50 on the second.

Now consider concurrent interval schedules, such as VI 30-sec reinforcement for pecks on one key and VI 60-sec reinforcement for pecks on the other. In this case, the pigeon could produce 120 reinforcers per hour by pecking only the first key, and 60 reinforcers per hour by pecking only the second key. By pecking both keys, however, it could produce all the reinforcers of both schedules, or 180

reinforcers per hour. For this reason organisms tend to distribute their responses to each of two concurrent VI schedules. But there is a complication. If a peck on one key is immediately followed by a reinforced peck on the other key, the reinforcer may affect the likelihood of the sequence of pecks. Thus, pecks on each key may come partially under the control of the reinforcers scheduled for the other key. To maintain the independence of the schedules, therefore, concurrent VI procedures often incorporate a *changeover delay*, which prevents either response from being reinforced immediately after a changeover from the other response (Herrnstein, 1961a; Catania and Cutts, 1963). With a changeover delay, the distribution of an organism's responses to two concurrent schedules tends to be proportional to the distribution of reinforcers arranged by the schedules (e.g., Herrnstein, 1961a; Herrnstein and Loveland, 1975). In the preceding case, the pigeon emits about twice as many pecks on the VI 30-sec key as on the VI 60-sec key.

This phenomenon is sufficiently general that it has been proposed as a general law of behavior, called the *matching* law (Herrnstein, 1970; see also Catania, 1976; Rachlin, 1971). The law states that the relative frequency of a response will match the relative frequency of the reinforcers produced by that response. The law even holds for concurrent ratio schedules, because exclusive responding on one schedule means that all the reinforcers will be delivered according to that schedule. Herrnstein's account has also been applied to the responding maintained by a single reinforcement schedule (e.g., Figure 8-2), on the assumption that other events besides the reinforcers arranged by the experimenter may have reinforcing effects, even though we cannot identify them.

The matching law effectively summarizes performances in a variety of schedules, but its empirical foundations have been controversial. For example, consider how concurrent VI schedules operate when arranged for a pigeon's pecks on two keys. As the pigeon continues to peck one key, more and more time passes during which the VI schedule for the other key may have made a reinforcer available. Thus, a time will be reached when the probability of reinforcement for changing over to the other key will be greater than the probability of reinforcement for continuing to peck the same key. It has therefore been argued that the pigeon distributes its responding to both keys in concurrent VI schedules because the pigeon always emits that response with the highest current probability of reinforcement, and this response shifts from one key to the other as time passes (Shimp, 1966). This performance has been described in terms of the principle of *maximizing*, which states that an organism with several responses available will emit the one with the maximum probability of reinforcement. In the case of unequal concurrent VR or FR schedules, this response is always the one with the smaller ratio, but in the case of concurrent VI schedules the response with the maximum probability of reinforcement changes from moment to moment; concurrent VI performance has therefore been called *momentary maximizing*.

Matching and maximizing seem to be contradictory alternatives. But they

are measured in different ways. We cannot speak about matching unless we have some sample of responses and reinforcers from which to estimate relative frequencies, but we can speak about maximizing with single responses, simply by noting whether the response was the one with the maximum probability of reinforcement. To some extent, the issues are concerned more with the level of detail at which performances are analyzed than with the performances generated by concurrent schedules (cf. the distinction between molar and molecular theory: chapter 5).

One feature of concurrent performances is that increases in the reinforcement of one response reduce the rate of other responses (e.g., Catania, 1963b, 1969a, 1973d; Rachlin and Baum, 1972; Wendt, 1936; B. A. Williams, 1975). This relation is implicit in some phenomena already noted. Consider the relation between response rate and rate of VI reinforcement (Figure 8-2). If the response rate




Figure 8-13. Derivation of the reduction in the rate of one response (A) by reinforcement of a concurrent response (B) from the VI rate function (curved line; cf. Figure 8-2) and the matching principle. The bracket shows the reduction in A produced by the reinforcement of B.

generated by a given rate of VI reinforcement is independent of how these reinforcers are distributed to two keys, it follows that an increase in the reinforcement of one response will decrease the rate of the other response. The relation is illustrated in Figure 8-13. The curved line shows an idealized version of the rate function for VI schedules. The bar marked A shows the rate of responding when response A alone is reinforced. If an equal VI schedule is arranged for response B, then the overall rate of reinforcement will double (x to $2x$). According to the matching principle, the total responding will now be equally distributed to the two

Variables that have small effects in single-response schedules sometimes have large effects in concurrent schedules. Concurrent schedules are therefore often used to study reinforcement variables (e.g., reinforcement duration: Catania, 1963a; delay of reinforcement: Chung, 1965; Chung and Herrnstein, 1967). Because concurrent schedules can make different consequences simultaneously available to an organism, these schedules provide appropriate baseline procedures for the study of choice or preference. A common arrangement is the concurrent-chain procedure illustrated in Figure 8-14 (Autor, 1960; Herrnstein, 1964b). Concurrent chains have shown that rate of reinforcement is a more important determinant of preference than the number of responses per reinforcement (Herrnstein, 1964b: see also Neuringer and Schneider, 1968; Rachlin, 1967). Among the preferences demonstrated with pigeons in concurrent chains are not only those among various parameters of reinforcement schedules (e.g., Fantino, 1977) but also preference for variable over fixed schedules (Herrnstein, 1964a), preference for informative stimuli over uninformative stimuli (Catania, 1975; Green and Rachlin, 1977), and preference for situations that allow free choices among alternative responses over situations that do not allow such choices (Catania, 1975: cf. the discussion of freedom in chapter 1). Concurrent-chain procedures that allow an organism to choose between situations differing in whether the organism can emit a response that commits it to a particular later alternative or does not permit such a commitment are pertinent to complex behavioral processes (Rachlin and Green, 1972). Analogous human behavior is sometimes discussed in the language of self-control, as when we say that someone makes a commitment to behave in a certain way or else yields to temptation (for a detailed account, see Rachlin, 1974).

SUMMARY: SCHEDULE ANALYSIS

The analysis of reinforcement schedules is relevant to various behavioral phenomena, ranging from causal relations between behavior and environment in simple schedules to such complex processes as self-control and freedom of choice. Yet even so-called simple schedules are not simple, and the complexity of schedule effects has made the analysis of schedules highly technical. We examined properties of variable and fixed ratio and interval schedules, and briefly sampled other major topics, including multiple, chained, and concurrent schedules. We often omitted procedural details. This is perhaps inevitable. Unlike most other areas in this text, reinforcement schedules did not exist systematically as a subject matter until relatively recent times (Skinner, 1938, 1956; Ferster and Skinner, 1957) and a major concern of this subject matter has typically been maintenance of behavior in the steady state (yet the change in behavior that accompanies any transition from one schedule to another is an instance of learning). The analysis of schedules forces us to recognize that response rate itself is a differentiable property of behavior along with topography, force, duration, and so on. Thus, an integration of

the diverse phenomena of reinforcement schedules is likely to come not from such specific measures as response rate and its derivatives, but rather from more general properties of behavior independent of response rate, such as the resistance of behavior to change (Nevin, 1974).

Respondent Behavior: Conditioning

CONDITIONAL REFLEXES

Types of conditioning

Conditioning and contiguity

Contiguity and consequences

Stimulus combinations in conditioning

Sensory preconditioning and second-order conditioning

Stimulus compounds

Overshadowing and blocking

Inhibitory stimuli in compounds

Autoshaping and automaintenance

Summary

OPERANT-RESPONDENT INTERACTIONS

Conditioning and emotion

Preaversive and preappetitive stimuli

RECAPITULATION: A BEHAVIOR PARADIGM

Kinds of contingent stimuli

Kinds of contingencies

Classification of contingencies and contingent stimuli

Relations among responses

Conclusion

NINE

The language of emotion includes a number of etymological clusters. For example, *eager*, *anger* and *anxiety* share the same roots, as do *choleric*, *melancholy*, *glad* and *glee*, and *wrath* and *worry*, and *sad* and *satisfy* (the last pair is also related to *satiate*). *Fear*, from the Indo-European *per-*, to try, to risk, to press forward or lead, has an extensive group of relatives that includes *experiment* and *apparatus*, *probability* and *opportunity*, *approach* and *deprivation*, and two contemporary synonyms for behavior, *comportment* and *performance*.

In respondent conditioning, we encounter a topic that once dominated the psychology of learning so much that it provided the opening chapters of most learning texts. In addition to *respondent conditioning*, it has gone by such names as *classical conditioning* and *Pavlovian conditioning*, and the language of *conditioned reflexes* has to some extent entered the everyday vocabulary (although in popular usage it has often been confused with instances of learning in operant behavior). The term *conditioned*, from the Russian условный, might better have been translated as *conditional*, because the name was applied to reflexes conditional upon relations among environmental stimuli. Respondent conditioning is an instance of stimulus control applied to stimulus-presentation operations rather than to the contingencies in consequential operations. In other words, instead of signalling the consequences of responding, a stimulus may simply signal the presentation of some other stimulus. Pavlov's conditioned salivary reflexes are the prototype example: when a bell or a light repeatedly signalled food in the mouth of a hungry dog, salivation began to be elicited by the signalling stimulus as well as by the food itself.

In discussing consequential operations, we had to speak of classes of responses rather than individual instances, because individual responses are never exactly repeated. We called these classes operants. Similar problems exist for elicited behavior. For example, successive elicitations of salivation by food may

differ in latency, quantity, viscosity, and other properties. It is therefore appropriate to extend the language of classes to responses defined in terms of stimuli that produce them. These classes are called respondents, and correspond to the behavior earlier called elicited or reflexive. Thus, salivation produced by food in the mouth is a respondent class; it must be distinguished from salivation produced by acid in the mouth, which is a different respondent class, and from spontaneous salivation, which is not a respondent class at all because no eliciting stimulus can be identified (we say that spontaneous salivation is emitted rather than elicited; if we could identify an eliciting stimulus, we would not call such salivation spontaneous). In this chapter, we first discuss the behavior generated by various types of Pavlovian or respondent conditioning. Then we consider how such behavior may interact with operant responding. We will find such interactions relevant to the topic of emotion.

CONDITIONAL REFLEXES

We produce respondent behavior by presenting stimuli, and we change respondent behavior by changing these stimuli. For example, by varying the concentration of acid in the mouth (e.g., dilutions of vinegar), we change the quantity of saliva elicited. This means we are limited in the extent to which we can modify respondent behavior. We can create new operants by shaping, but because the properties of respondents are determined by their eliciting stimuli, no equivalent procedure exists for respondent behavior. We can, however, alter which stimuli will elicit particular responses. Let us examine Pavlov's (1927) conditioning procedures in more detail.

We begin with a dog in a harness. One of its salivary ducts is connected to a system that allows the dog's salivation to be recorded. Two stimuli are used: the sound of a bell beside the dog, and food in the form of pellets or powder that can be delivered directly into the dog's mouth. First we examine the effects of each of these stimuli when presented separately. When the bell first rings, the dog pricks up its ears and turns its head toward it (these have sometimes been called *orienting* responses). With successive rings, however, these responses diminish in magnitude, perhaps even becoming undetectable; they can be reinstated by waiting some time before ringing the bell again (cf. adaptation, chapter 2). When food is placed in its mouth, the dog chews and salivates. Although these responses diminish somewhat over successive food presentations, their magnitude remains substantial over the course of an experimental session.

Suppose now that the bell is used to signal food. If the bell rings five seconds before each food delivery, we may find after a number of combined presentations that salivation sometimes begins in the five-second period between bell and food and that the bell is sometimes followed by salivation even on an occasional trial when we deliberately omit the food. In neither case can we attribute the salivation to food as an eliciting stimulus: in the first case the salivation began before the food was presented, and in the second the food was not presented at all. To the

extent that the bell is now a stimulus that has acquired the power to elicit salivation, we say that we have created a new respondent class, salivation elicited by the sound of the bell. We call the relation between bell and salivation a conditional reflex because it is conditional on the prior relation of bell and food.

The sequence of events is illustrated in Figure 9-1. The bell at first elicits orienting responses, but these responses disappear with repeated presentations; at this point, the bell may be called a *neutral stimulus*. In an unconditional reflex, food elicits salivation; in this relation, the food is called an *unconditional stimulus* or *US* and the salivation is called an *unconditional response* or *UR*. Conditioning begins when the bell reliably precedes food; at this point, the bell still has no effect on salivation, and we may continue to regard it as a neutral stimulus. After a period of conditioning, a *conditional reflex* may be established; the bell may elicit salivation before food is presented (*a*), or the bell may elicit salivation even when food is omitted on an occasional trial (*b*). The bell may now be called a *conditional stimulus* or *CS* and salivation elicited by the bell a *conditional response* or *CR*.

The difference between a conditional stimulus and an unconditional stimulus is not simply which comes first. If we reversed their order, we would not expect the eliciting effects of food followed by bell to be much different from food alone. In fact, the relative effectiveness of stimuli as conditional and unconditional stimuli may be predictable from the probabilities with which the stimuli elicit their respective responses. For example, consider salivation elicited by food in a dog's mouth and leg flexion produced by shock to the dog's leg. A mild shock may become a conditional stimulus eliciting salivation if it is reliably followed by food, but a




Figure 9-1. Relations between stimuli and responses in respondent conditioning. An initially neutral stimulus (BELL: NS) is followed by an unconditional stimulus (FOOD: US) that elicits salivation. If the neutral stimulus begins to elicit responding like that elicited by the unconditional stimulus, the neutral stimulus is called a conditional stimulus (CS). OR - orienting response; UR - unconditional response; CR - conditional response.

strong shock followed by food is not likely to do so. On the other hand, food may become a conditional stimulus eliciting leg flexion if it is reliably followed by a strong shock, but food followed by a weak shock is not likely to do so. The differences are consistent with the relative probabilities with which food may elicit salivation and with which mild or strong electric shock may elicit leg flexion.

Types of conditioning

The temporal relations between the two stimuli can be arranged in varied ways. Situations in which the onset of the conditional stimulus precedes the onset of the unconditional stimulus by no more than five seconds are sometimes arbitrarily grouped together as instances of *simultaneous conditioning*. This usage is perhaps based on two circumstances; that the optimal interval between conditional and unconditional stimuli is roughly half a second (e.g., B. Reynolds, 1945; Kimble, 1947), and that omissions of the unconditional stimulus sometimes reduce the reliability with which conditional responding is maintained. Short intervals like half a second provide little opportunity to observe conditional responding. The choice then was to omit the unconditional stimulus on some trials, or to lengthen the interval between stimulus onsets. Because of the effects of stimulus omissions on the reliability of conditioning, lengthening the interval between stimulus onsets became favored over occasionally omitting the unconditional stimulus, and the distinction between strict simultaneity and these relatively short delays was overlooked.

The effect of occasional omissions of the unconditional stimulus had theoretical ramifications. In a usage that is becoming rare, presentations of the unconditional stimulus in respondent conditioning were called reinforcement, and therefore the procedure in which this stimulus was omitted on occasional trials was called partial reinforcement. Distinctions between operant and respondent conditioning were then argued on the basis of the so-called partial reinforcement effect or PRE: relative to reinforcement of every response, partial reinforcement generated substantial quantities of responding in operant behavior, as we saw in chapter 8, but the analogous procedure in respondent conditioning seemed to reduce responding. Now that the term reinforcement has become more restricted in its scope, the comparison no longer seems relevant. In its time, however, it provided one of the most persuasive grounds for distinguishing between the two types of conditioning.

Some arrangements of conditional and unconditional stimuli are contrasted with simultaneous conditioning in Figure 9-2. In both *trace conditioning* and *delay conditioning* a relatively long time elapses between the onset of the conditional stimulus and the onset of the unconditional stimulus. They are distinguished by whether the conditional stimulus turns off or remains present during that time. In both trace and delay conditioning, conditional responding at first occurs shortly after the onset of the conditional stimulus, but with successive trials it gradually moves closer to the time at which the unconditional stimulus is to be




Figure 9-2. Schematic presentation of various respondent-conditioning procedures, using a bell as an example of a conditional stimulus and food as an example of an unconditional stimulus. In simultaneous conditioning, bell is followed within less than five seconds by food. Different time relations are illustrated in trace, delay, and temporal conditioning, and the stimulus order is reversed in backward conditioning. In differential conditioning, a bell is followed by food but a tone is not. CS - conditional stimulus; US - unconditional stimulus; NS - neutral stimulus.

delivered. Trace conditioning acquired its name from the assumption that the conditional stimulus had to leave some trace in the organism's nervous system for such conditioning to be effective. But successive presentations of the unconditional stimulus itself at regular intervals (e.g., every half-hour) also produce conditional

responding; this procedure is called temporal conditioning, and responding has sometimes been spoken of as conditioned to time as a stimulus (temporal conditioning involves repeated presentation of a stimulus, and so is equivalent to some of the stimulus-presentation operations in chapter 3). When the order of the two stimuli is reversed, the procedure is called *backward conditioning*. Because this procedure is not particularly effective in producing conditional responding, it is perhaps inappropriate to speak of a conditional stimulus (the stimulus is labelled neutral in Figure 9-2). Finally, a procedure in which one stimulus is established as a conditional stimulus through its relation to the unconditional stimulus while a second stimulus is not so established because it never precedes the unconditional stimulus is called *differential conditioning*. The stimuli are sometimes called positive and negative conditional stimuli (CS+ and CS-).

Conditioning and contiguity

Some of the attention historically given to respondent conditioning may have depended on how easily it could be related to the concept of association, a principle of learning with substantial precedent in philosophy and psychology. Learning had been said to take place through the association of ideas, and conditional reflexes seemed to represent a primitive example of the formation of such associations. If ideas were associated, it was argued, then the occurrence of one would lead to the other. In a kind of mental chemistry, ideas were supposed to become associated through such properties as having common elements or occurring together in time. It remained then to suggest that ideas could be interpreted as responses generated by environmental events, so that if events occurred together in the past remembering one would call up the other. We need not dwell on the details. Even those parts of the contemporary psychology of learning that can still be called associationistic have evolved considerably from earlier formulations. We will see that respondent conditioning cannot be interpreted simply as making one stimulus a substitute for another. For the present, the point is that respondent conditioning was regarded as a process at the root of all learning, and was assumed to take place merely through the temporal contiguity of events, their occurrence together in time. Theoretical debates then revolved around the primacy of respondent conditioning and other processes, and particular attention was given to finding ways of interpreting instrumental or operant behavior as an instance of behavior generated by respondent principles (for various sides of the argument, see Guthrie, 1935; Hull, 1943; Konorski, 1948; Mowrer, 1960; Schlosberg, 1937; Skinner, 1935; K. Smith, 1954).

Part of the problem lies in the definition of contiguity. Contiguity was often interpreted in terms of the number of times the stimuli had occurred together, or the number of stimulus pairings. Even when the number of pairings remains constant, however, the relation between conditional and unconditional stimuli can vary. Assume that a bell (S1) and food (S2) are arranged within trials, and that we can ignore the stimuli that demarcate the trials. The rows in Figure 9-3 show




Figure 9.3. Conditional relations between two stimuli, $S1$ and $S2$. Each row represents a sample of ten trials from a conditioning procedure. Each case involves exactly three pairings of $S1$ and $S2$, in trials 2, 5 and 7, but only in the first case is $S1$ positively correlated with $S2$. In the second case $S2$ is as likely given $S1$ as not given $S1$, and in the third case the likelihood of $S2$ is less given $S1$ than not given $S1$. Only in the first case will $S1$ be established as a conditional stimulus relative to $S2$ as an unconditional stimulus. The probabilities of $S2$ given $S1$ and of $S2$ given no $S1$ are shown at the right. (Cf. Rescorla, 1967.)

samples of ten trials from three conditioning procedures. In each, $S1$ is paired with $S2$ in trials 2, 5, and 7. In the top row, only those trials include $S2$, which is therefore perfectly correlated with $S1$: the probability of $S2$ given $S1$ is 1.0, but it is zero without $S1$. In the middle row, $S2$ occurs in every trial, and therefore $S1$ is irrelevant to whether $S2$ occurs: the probability of $S2$ is 1.0 given $S1$ or given no $S1$. In the bottom row, $S1$ occurs in six trials, but in only half of these is it followed by $S2$, whereas $S2$ occurs in three-quarters of the trials in which $S1$ is not presented: the probability of $S2$ given $S1$ (0.50) is lower than the probability of $S2$ given no $S1$ (0.75). Only in the first procedure, $S1$ might even reduce the likelihood of conditional responding elicited by the trial stimulus. The conditional relation between the two stimuli rather than the number of pairings is the appropriate basis for classifying conditioning procedures¹ (Rescorla, 1967).

Contiguity and consequences

Many conditional responses have been established through respondent procedures (e.g., see Hull, 1934). Pavlov's salivary conditioning is probably the most familiar, but other studies demonstrated conditioning based on such unconditional reflex relations as the knee jerk elicited by a tap of the patellar tendon (Twitmyer, 1902) and limb withdrawal elicited by electric shock (Bechterev, 1933).

¹In fact, if $S1$ had elicited salivation even when it was explicitly unpaired with food, we would speak of *pseudoconditioning*. For example, a startle response might be elicited by a variety of innocuous stimuli after a traumatic shock, even though these stimuli had never occurred together with the shock; thus, pseudoconditioning control procedures are sometimes incorporated into conditioning procedures, to make sure that observed effects depend on the relation between the conditional and unconditional stimuli rather than simply on the presentations of the unconditional stimuli.

The conditioning of eye blinks elicited by a puff of air to the eye provided substantial data (e.g., Gormezano, 1972), but the conditioning of pupillary constriction elicited by light to the eye, once believed feasible, has not been successful (Young, 1958). The Russian literature includes a variety of conditioning demonstrations (e.g., Bykov, 1957). For example, for several days a dog was placed in a waiting area at neutral temperature before being moved into a heated room. The decrease in metabolism and oxygen consumption that typically occurred in the heated room began to occur in the waiting area. On the other hand, after several days in which the stay in the waiting area was followed by being moved into a cool room, the increased metabolism and oxygen consumption that typically occurred in the cool room began to occur in the waiting area (Bykov, 1957, pp. 183-210).

Those instances of conditioning based on aversive stimuli such as electric shock were called *defensive conditioning*, on the assumption that the responses elicited by such stimuli occurred because they had some natural defensive functions (cf. chapter 5 on species-specific defense reactions; Bolles, 1970). One frequently cited example is an experiment with an infant boy named Albert (Watson and Rayner, 1920). The aversive stimulus was the sound, just behind Albert, of striking a suspended steel bar with a hammer. This sound produced startle or crying or withdrawal responses. When the sound followed presentations of a white rat, these responses began to occur in the presence of the rat and of other stimuli having properties in common with the rat, such as cotton wool. Watson and Rayner called these responses *conditioned emotional reactions*. Yet the details of the experimental procedure show that the initial clangs of the steel bar were not independent of behavior:

1. White rat suddenly taken from the basket and presented to Albert. He began to reach for the rat with left hand. Just as his hand touched the animal the bar was struck immediately behind his head. The infant jumped violently and fell forward, burying his face in the mattress. He did not cry, however.
2. Just as the right hand touched the rat the bar was again struck. Again the infant jumped violently, fell forward and began to whimper. In order not to disturb the child too seriously no further tests were given for one week. (Watson and Rayner, 1920, p. 4)

At least at the outset, therefore, Watson and Rayner's experiment used punishment of reaching toward the rat and not merely response-independent stimulus presentations. This was presumably important in getting Albert to attend to the white rat as a discriminative stimulus. Nonetheless, we cannot unambiguously attribute little Albert's responses to conditioning. We cannot even rule out the consequences of Albert's responses to the struck bar; after all, these responses led the experimenters to terminate the procedure for a week.

The problem, however, is not restricted to Watson and Rayner. Once we become alert to the possibility of response consequences in supposed conditioning procedures, we can often find such consequences. For example, in early demonstrations of limb withdrawal elicited by shock, the method of attaching the electrodes was often a matter of indifference. Yet if both electrodes were attached to a dog's

leg, a flexion could not prevent shock delivery, whereas if either or both electrodes were attached to the floor on which the dog's paw rested, a flexion prevented or terminated shock delivery because once lifted from the electrodes the paw no longer completed the electrical circuit. In fact, the classic defensive reflex of Bechterev (1933) typically had both electrodes on a surface the organism touched, so that a response prevented or terminated the shock; Bechterev was therefore probably studying avoidance and escape behavior rather than respondent conditioning. Recognizing the different implications of the two methods of electrode placement for the consequences of responding was an important step in the evolution of the distinction between operant and respondent behavior (cf. Schlosberg, 1937; Skinner, 1935; see also Kimmel, 1976).

Once consequences became implicated in a few cases of presumed respondent conditioning, it was tempting to seek them in all cases. For example, even if both electrodes are attached to a dog's leg in defensive conditioning, might not the flexion have consequences? Suppose that a bell reliably precedes electric shock. How can we tell whether the shock is as aversive when it passes through a flexed limb? In this instance of respondent conditioning, the conditional response is ordinarily slower and different in magnitude than the unconditional response. Furthermore, if the dog's leg is already flexed when shock is delivered, less postural adjustment is required than if the dog stands on all fours and then shifts its weight to the remaining three legs when flexing the shocked leg (Wagner, Thomas, and Norton, 1967). Clearly salivation too has consequences; it affects taste and swallowing in the case of dry food and dilution of the solution in the case of acid on the tongue.

The introduction of respondent conditioning into learning theory began with attempts to reduce all instances of operant learning to special cases of respondent conditioning, but with such arguments the situation had been turned around. The case was made that all instances of respondent conditioning could be interpreted in terms of response consequences that were earlier unnoticed. The next step was to observe that autonomic responses such as salivation and constriction or dilation of blood vessels were often accompanied by somatic responses (i.e., action of the striate muscles, which produce skeletal movement). The position could then be taken that autonomic responses in respondent conditioning were artifacts, in that they incidentally accompanied behavior established through operant or instrumental processes (K. Smith, 1954). The status of respondent conditioning therefore came to depend on demonstrations of conditioning that could not be interpreted in terms of the consequences of responding.

One approach was to see whether conditional responding could be modified by explicitly arranged consequences. If such consequences did not modify the conditional reflex, then the argument that the new reflex relation depended on other unidentified consequences would no longer be particularly convincing. Sheffield (1965) therefore added consequences to the conditional salivation generated by the classical Pavlovian situation. Specifically, a conditional stimulus preceded food, but food was omitted if the dog salivated on that trial. (The procedure is an example of negative punishment sometimes referred to as *omission train-*

ing). In other words, this arrangement converted the standard Pavlovian procedure to one in which the consequence of salivating was no food and the consequence of not salivating was food. Sheffield concluded that salivation did not come under the control of its consequence, the omission of food. One basis for this conclusion was the performance of the dog Vicki. At the beginning of training, conditional salivation had not yet been established, and therefore the conditional stimulus, a tone, was consistently followed by food. These circumstances produced conditional salivation, but in trials on which Vicki salivated, food was omitted; Vicki's conditional salivation therefore decreased. Once conditional salivation decreased, the tone was again consistently followed by food, and Vicki's conditional salivation reappeared. The cycle of conditional salivation, omitted food, decreased salivation, reinstatement of food, and return to conditional salivation was repeated, and for forty days (800 trials) Vicki's cyclic conditional salivation remained fairly stable. Although Vicki could have received food on every trial by not salivating during the tone, Vicki did not learn to do so and received food on only some of the trials each day.

It would be premature to conclude that this settled the issue. A consequence effective as a reinforcer for one response may not be effective for some other response (chapter 4). Given that food elicits salivation, we should not be surprised if not salivating is ineffectively reinforced by food. Later research in fact demonstrated that salivation could be reduced when its consequence was the omission of a reinforcer. In that case, however, the reinforcer was water, which does not itself elicit salivation (N. E. Miller and Carmona, 1967). Salivation, elicited in some circumstances, can be modified by its consequences under others. Our task, for this response and others, is to determine the conditions under which either of these processes can occur.

Stimulus combinations in conditioning

We have considered cases of respondent conditioning based on unconditional stimuli that serve in other situations as reinforcers (e.g., food) or punishers (e.g., shock). Their advantage for studies of conditioning is that they have substantial eliciting effects. We could also ask about conditioning when relatively neutral stimuli such as lights or sounds serve as unconditional stimuli. The problem is that conditioning is difficult to assess when no elicited responses can easily be measured. Two procedures concerned with the effects of such stimuli, called *sensory preconditioning* and *second-order conditioning*, are illustrated for conditional leg flexions in dogs in Figure 9-4.

Sensory preconditioning and second-order conditioning. First consider sensory preconditioning (e.g., Brogden, 1939). In the first phase of this procedure, preconditioning, one stimulus signals a second stimulus. In the example of Figure 9-4, a bell is consistently followed by a tone. In the second phase, a conditional reflex is established in which the second stimulus becomes a conditional stimulus.




Figure 9-4. Schematic diagram of the phases of sensory preconditioning and second-order conditioning procedures, using bell and tone as conditional stimuli and shock-elicited leg flexions in a dog as the unconditional response.

In the figure, tone is followed by shock. Once the conditional reflex is established, so that tone elicits leg flexion, the eliciting effects of the bell are tested. Leg flexion to the bell is then taken to mean that the bell became a conditional stimulus relative to the tone during preconditioning.

In second-order conditioning, the order of phases is reversed. In the example in Figure 9-4, a conditional reflex in which tone elicits leg flexion is established first, by presenting tone and then shock. Later, bell is followed by tone. In this instance, the question is whether the conditional stimulus established in the first phase can function as a unconditional stimulus for another stimulus in the second phase. The difficulty is that the tone may lose its effectiveness as a conditional stimulus as it is repeatedly presented without shock in the second phase, but during this time the bell must acquire its conditional properties. An alternative procedure, in which bell is followed by tone and then by shock on all trials within a single phase, would be ambiguous; if the bell did come to elicit leg flexions we would not know whether these occurred because of the relation between bell and tone or because of the relation between bell and shock.

If the conditional stimulus is presented alone after sensory preconditioning, thereby extinguishing the conditional reflex, the preconditioning stimulus also will no longer elicit a conditional response. In other words, in the example of Figure 9-4, presenting tone alone after the conditioning phase until the tone no longer elicits flexions will also make the bell lose its effectiveness in eliciting flexions. But the comparable procedure after second-order conditional reflex (Rizley and Rescorla, 1972). In the example of Figure 9-4, presenting tone alone until it no longer elicits flexions will not eliminate the conditional flexions to the bell established during the second-order phase. This outcome is paradoxical. Consider an imaginary human case history. A young man sees blood in painful circumstances, and the sight of blood therefore becomes a conditional stimulus that elicits those emotional responses we call fear. At a later time (equivalent to the second-order phase), the young man sees blood in an elevator, and thereby acquires a fear of

elevators. The young man later goes into hospital work, and in that context gradually gets over his fear of the sight of blood. According to Rizley and Rescorla's findings, this change will not reduce his fear of elevators, even though the conditional response to the blood had been the basis for that fear. This case is imaginary, however, and we must be cautious in the generalizations we draw from experimental situations. In any event, these findings make it clear that respondent conditioning cannot be interpreted as simply establishing one stimulus as a substitute for another.

Stimulus compounds

Whether a stimulus becomes effective as a conditional stimulus depends on the context of stimuli within which it appears (e.g., Kamin, 1969). Suppose, for example, that loud tone and dim light are presented together and followed by some unconditional stimulus, such as a shock that elicits leg flexion. Once conditional responding developed to this stimulus pair, we could present each stimulus separately. We might find tone much more effective as a conditional stimulus than light. (The concept of attention, as in operant discrimination, is relevant to respondent conditioning; we might say that the organism was attending more to tone than to light, or that tone was more salient than light.) The contributions of the separate stimuli that make up a stimulus compound have been the subject of extensive mathematical analyses (e.g., Rescorla and Wagner, 1972), but we will restrict our examples to some basic effects of stimulus compounding.

Overshadowing and blocking. The tone-plus-light example assumed that the organism had no history of conditioning to either stimulus of the stimulus compound. The procedure is shown schematically as *overshadowing* in Figure 9-5. When the stimuli of a compound do not become equally effective conditional stimuli, we say that the more effective stimulus overshadowed the less effective stimulus. But such an effect may also occur when one of the two stimuli already has a history of conditioning. For example, suppose dim light had been established as a conditional stimulus by itself before it was presented together with tone. If simultaneous light and tone were then followed by shock until conditional flexions were observed, we might find when presenting the stimuli separately that light but not tone was effective as a conditional stimulus, even though tone preceded the shock as reliably as did light. When a stimulus fails to become effective as a conditional stimulus because it is presented together with some other already effective stimulus, we say that the stimulus with the prior history blocked conditioning to the new stimulus. This procedure is shown schematically as *blocking* in Figure 9-5.

Analogous phenomena occur in operant discriminations, and remind us that operant discrimination and respondent conditioning are both instances of stimulus control. In operant discrimination, a stimulus signals the consequences of responding, and the behavior generated by those consequences tends to occur more often when the stimulus is present than when it is absent. In respondent condition-




Figure 9-5. Schematic illustrations of overshadowing and blocking. In overshadowing, a conditioning procedure in which neither of the two stimuli of a compound had a prior history of conditioning establishes only one of the two stimuli as a conditional stimulus (or, alternatively, makes one of the stimuli more effective than the other). In blocking, conditioning is established with respect to one stimulus, and this history prevents the other stimulus from becoming effective as a conditional stimulus when the stimuli are presented as a compound (or, alternatively, reduces the effectiveness of the second stimulus). S1 and S2 - stimuli; US - unconditional stimulus; CR - conditional response.

ing, one stimulus signals the presentation of another stimulus, and the behavior generated by the second stimulus similarly tends to occur more often when the first stimulus is present than when it is absent. In each case, whether responding is consequential or elicited, it comes to be concentrated in the presence of the discriminative stimuli.

Inhibitory stimuli in compounds. Stimuli can signal the omission as well as the presentation of other stimuli. When they do so, they sometimes acquire the capacity to reduce the effectiveness of conditional stimuli, and are described as *inhibitory*. These effects can be measured most directly with stimuli presented in combination. An example is provided in Figure 9-6, which illustrates a conditioning procedure involving food-elicited salivation in a dog. First, a bell is established as a conditional stimulus (CS_A^+). Once the bell comes reliably to elicit salivation, a tone is presented either by itself or together with a light on irregularly alternating trials. When tone (CS_B^+) is presented alone, it is followed by food. When it is presented with the light (CS_X^+), the food is omitted. Eventually tone alone elicits conditional salivation, but the pairing of tone and light does not. It might be assumed that the dog simply discriminates tone alone from the tone-plus-light combination. But the inhibitory effect of the light is demonstrated when later pre-




Figure 9-6. Schematic illustration of a procedure for demonstrating an inhibitory component of a stimulus compound. First, the succession of bell and food establishes conditional salivation to the bell. Next, a tone alone is followed by salivation is established to tone alone but not to tone plus light. When the light is presented together with the bell, the bell elicits less salivation than it would if presented alone.

sentations of the bell together with the light elicit less conditional salivation than the bell would have elicited if presented by itself.²

Autoshaping and automaintenance

The cases of respondent conditioning so far have included both autonomic responses (e.g., salivation) and somatic or skeletal responses (e.g., leg flexion). Chapter 3 considered how these two classes contributed to theoretical distinctions between operant and respondent behavior. The demonstration that autonomic responses such as salivation or intestinal motility could be modified by their consequences had substantial impact on these theories. That impact was paralleled by the demonstration that somatic or skeletal responses could be affected by respondent procedures. Both demonstrations implied that operant and respondent processes could not be distinguished by kinds of responses; the critical difference was instead in the respective consequential and stimulus-presentation operations.

We discussed earlier the ambiguity of experiments on conditional leg flexions to shock. The problem was that procedures could not be designed to guarantee that leg flexions would have no consequences. Perhaps partly for this reason, the demon-

²In another procedure, we could follow the tone-light combination but neither of the stimuli alone by the unconditional stimulus. In this case, we would eventually find that the combination but not the individual stimuli would elicit the conditional response, but we could not easily appeal to either stimulus as inhibitory. This demonstration simply shows that organisms can respond differentially to relations among stimuli as well as to the individual stimuli themselves.

stration of respondent conditioning of the pigeon's key peck, in a procedure called *autoshaping* (P. L. Brown and Jenkins, 1968; Hearst and Jenkins, 1974; Schwartz and Gamzu, 1977), received special attention. In addition, because the key peck was a commonly chosen response for the study of consequential responding, it was important to determine the extent to which respondent processes might enter into those performances.

The autoshaping procedure was introduced as a convenient alternative to shaping the key peck by successive approximations (chapter 6). Autoshaping begins with a pigeon that eats reliably from the feeder but has not pecked the key in a standard chamber. At irregular intervals the key is lit, and several seconds later the feeder is operated independently of the pigeon's behavior. Thus, the lit key becomes a discriminative stimulus (or conditional stimulus) that signals another stimulus, food. The food occasions eating (and in the pigeon eating includes pecking). We therefore might say that food is an unconditional stimulus and that pecking food is an unconditional response. After a few presentations of lit key and then feeder, the pigeon begins to face and move toward the key when it lights. Within perhaps fewer than ten and rarely more than one hundred trials, the pigeon will begin to peck the key whenever it is lit. After pecking is established by autoshaping, the continuation of the procedure is called *automaintenance*. *Autoshaping* and *automaintenance* simply distinguish the changes in behavior leading up to the first peck from the maintained behavior following this first peck.

In autoshaping and automaintenance, food deliveries occur independently of behavior. It is therefore difficult to attribute autoshaped pecking to its consequences. Nevertheless, once such pecking begins it will often be followed by food. An argument based solely on the observation that no consequences of pecking are obvious may not be persuasive. Autoshaped pecks were therefore studied in an omission procedure analogous to Sheffield's experiment with salivation: food was delivered after the key light only on trials in which the pigeon did not peck the key (D. R. Williams and Williams, 1969). As with salivation, pecks were maintained in a substantial proportion of trials even though those pecks caused the omission of food. Key pecking persisted over hundreds of trials, the behavior presumably stabilizing at a level at which enough trials occurred without pecks (and therefore with food) to maintain pecking in the other trials.

When food is repeatedly presented to a hungry pigeon, pecking becomes a more and more dominant component of behavior as time passes between food presentations (cf. chapter 3 on interim and terminal behavior). Autoshaped key pecking in the pigeon may therefore be interpreted as terminal behavior under the stimulus control of the key light: pecks generated by repeated food presentations tend to occur mainly in the presence of the key light, and to be directed toward this stimulus so strongly that they strike the key and are recorded as key pecks. Once autoshaped pecking begins, it may be maintained indefinitely by repeated presentations of key light and food even though it has no obvious consequences (in fact, relative to keeping its head in or near the feeder, the pigeon's key pecking may delay its access to food).

The production and maintenance of key pecking in autoshaping and auto-

maintenance has the critical features that define respondent conditioning. The respondent vocabulary is therefore appropriate. We may speak of the key light as a conditional stimulus. It acquires its capacity to elicit a conditional response, key pecking, through its correlation with food. Food is therefore the unconditional stimulus, which elicits an unconditional response, pecking. As in other cases, the correlation between key light and food, not simply the number of pairings of key light and food, determines whether autoshaped pecking occurs (Bilbrey and Winokur, 1973; Gamzu and Williams, 1973). An additional and distinctive feature of autoshaping is the directed nature of the conditional response: pecks generated by the key light might occur anywhere in the chamber (in the air, at the wall, around the food hopper), but instead they are directed at the key. The inverse relation also holds: pigeons tend to move away from stimuli correlated with the absence of food (Wasserman, Franklin, and Hearst, 1974). For example, if a green key precedes food deliveries but a red key does not, autoshaped pecking occurs when green is presented, but the pigeon moves to the opposite side of the chamber when red is presented. The directed nature of autoshaped pecking has been called sign-tracking (Hearst and Jenkins, 1974). A response must be emitted before it can be reinforced (chapter 4), and it may sometimes be emitted because it was once elicited (chapter 3). Perhaps then the behavioral relations that occur in autoshaping are prototypes of the processes from which operant behavior evolves.

Chapter 8 considered behavioral contrast, the increase in the rate of responding in an unchanged component of a multiple schedule when the reinforcement rate in the other component decreased. It attributed that increase to the generation of autoshaped pecks by the differential correlation of the multiple-schedule stimuli with food. One type of evidence was the separation of key pecking into two separate classes on the basis of duration and topography (K. Keller, 1974; Schwartz and Williams, 1972). The topography of autoshaped or automaintained pecks is affected more by the nature of the unconditional stimulus than by contingencies (H. M. Jenkins and Moore, 1973). For example, the pigeon's pecking at grain is a briefer response than and differs in form from the pigeon's drinking of water, and when autoshaping produced by key light and food is compared with that produced by key light and water, the autoshaped pecks in each case resemble the responses elicited by the respective unconditional stimuli. When autoshaped pecks are based upon food, the pigeon looks as if it is eating the key, whereas when autoshaped pecks are based upon water, the pigeon looks as if it is drinking the key.

Summary

Our treatment of the basic components of respondent conditioning, conditional and unconditional stimuli and responses, was followed by a brief survey of types of conditioning, including simultaneous, trace, delay, temporal, and backward conditioning. We distinguished between pairing and predictability in defining contiguity in conditioning: even though one stimulus is paired with another stimulus, it will not become an effective conditional stimulus unless the second stimulus is more likely after the first stimulus than otherwise. We then examined

some empirical findings with various stimulus combinations in conditioning, including sensory preconditioning, second-order conditioning, overshadowing and blocking, and inhibitory stimuli in stimulus compounds. We concluded with an example of the conditioning of a skeletal response, the pigeon's peck, in autoshaping and automaintenance.

OPERANT-RESPONDENT INTERACTIONS

Behavior can sometimes be analyzed as a combination of operant and respondent processes. These processes can interact with each other when respondent-conditioning procedures are combined with operant or consequential procedures. For example, a stimulus that reliably precedes or signals an electric shock may itself come to elicit responses such as leg flexions, but it may also interrupt other behavior that has been maintained by its consequences, such as lever presses maintained by food reinforcement. We sometimes describe comparable behavior in humans in terms of fear or anxiety; procedures like these have often been regarded as relevant to emotion.

Conditioning and emotion

Stimuli that signal the presentation of other stimuli may be superimposed on baselines of ongoing operant behavior. For example, suppose that a rat's lever presses are maintained by an interval schedule of food reinforcement; from time to time a tone is presented, and at the end of the tone shock is delivered. In these circumstances, the tone typically reduces lever pressing, especially as the time approaches when shock will be delivered. This phenomenon, originally demonstrated by Estes and Skinner (1941), has been given various names: *anxiety*, *conditioned suppression*, and *conditioned emotional response* or *CER*. It is illustrated in Figure 9-7, which shows development of suppression and then recovery from suppression after shock was discontinued (Geller, 1960). The rat's lever presses were maintained by a VI 2-min schedule of food reinforcement, and three-minute presentations of tone were followed by shock. After the tone came to suppress responding, shock was discontinued, and lever pressing in tone recovered to earlier levels.

The procedure is an instance of respondent conditioning: one stimulus, tone, signals another stimulus, shock. In this instance it is not obvious that tone elicits responses similar to those elicited by shock. The tone suppresses reinforced lever pressing, but lever pressing begins again soon after shock has been delivered. If we looked more closely at the rat's behavior during the tone, we would find that many other classes of responses were affected besides lever pressing (i.e., heart rate, respiration: cf. Rescorla and Solomon, 1967; Blackman, 1977). We are most likely to invoke the language of emotion when an event has wide-spread effects across many different response classes, and we may therefore be tempted to speak of the

rat's fear or anxiety. But if we do so, we must recognize that such terms do not explain the rat's behavior. It would not do to say later that the rat stopped pressing the lever during tone because it was afraid; the effect of the tone on the rat's lever pressing led us to speak in terms of the rat's fear in the first place.

Our language of emotions is complex. We speak of emotions in others and in ourselves on the basis of both situations and the behavior that occurs in those situations (e.g., Bem, 1967). For example, we might speak of the behavior produced by preaversive stimuli in terms of fear or anxiety, but if we observed aggressive behavior directed toward another organism we would be more likely to




Figure 9-7. Development of and recovery from suppression in the presence of a preaversive stimulus. A rat's lever presses were maintained by a VI 2-min schedule of food reinforcement. Reinforcements are not shown on the cumulative records. The records on the left show the effects of superimposing a three-minute tone followed by electric shock on the baseline lever pressing. The dashed vertical lines mark off the periods of tone, during which the response pen was also displaced downward. By the 27th presentation, lever pressing was almost completely suppressed by the tone. The records on the right show recovery from suppression when the tone was no longer followed by shock. (Adapted from Geller, 1960, Figure 3.)

speak of anger. In either case, we must be clear that these are names for certain behavioral effects; they are not explanations. Consider, for example, the question of why someone acts in a particular way. The answer might be that the person is sad or depressed. If we then ask how we know that the person is sad or depressed, we perhaps get the answer that we can tell from how the person acts. We should not be satisfied with this. What more has been said than that the person is acting that way because the person is acting that way? It would be more useful to be told that the person is acting that way because of some specific event, such as the loss of a job or the breaking up of a love affair.

Preaversive and preappetitive stimuli

Although the language of emotion is important in our interactions with other people, it has not proved particularly useful in behavioral analyses of the effects of preaversive stimuli. Instead, the interactions between respondent conditioning and operant behavior, as when preaversive or preappetitive stimuli are superimposed on reinforced responding, have been analyzed more effectively in terms of various experimental parameters, such as baseline reinforcement schedule, baseline response rate, and so on.

The finding that positively reinforced responding may be suppressed by preaversive stimuli (Estes and Skinner, 1941) was later supplemented by the finding that negatively reinforced responding (avoidance: see chapter 5) may be enhanced by such stimuli (Sidman, Herrnstein, and Conrad, 1957). In other words, a rat whose lever presses avoid electric shocks may increase rather than decrease its lever pressing during a stimulus that precedes an unavoidable or inevitable shock. This enhanced responding has been given a variety of names, including conditional or conditioned enhancement or facilitation or acceleration. Once such enhanced responding develops during negatively reinforced responding, it may continue with responding maintained by positive reinforcement. For example, rhesus monkeys' lever pressing maintained by orange juice reinforcers was originally suppressed in the presence of a clicking noise that preceded shock, but after the monkeys acquired a history of lever pressing that avoided shock and were returned to the initial procedure, lever pressing during the preaversive stimulus was enhanced rather than suppressed (Herrnstein and Sidman, 1958).

The situations were then extended to superimposing preappetitive rather than preaversive stimuli on baseline operant behavior. For example, a key light that preceded response-independent food deliveries increased pigeons' key pecking when this stimulus was superimposed on key pecking maintained by DRL reinforcement (Herrnstein and Morse, 1957). Just as suppression during preaversive stimuli might be spoken of in terms of anxiety, it was tempting to speak of such enhancing effects of preappetitive stimuli in terms of joy. It was then assumed that the suppression of positively reinforced and enhancement of negatively reinforced responding by preaversive stimuli might be paralleled by the enhancement of positively reinforced and suppression of negatively reinforced responding by preappetitive stimuli (e.g., Leitenberg, 1966; Rescorla and Solomon, 1967). But, aside from the problems of distinguishing such phenomena from other effects of conditioning (e.g., autoshaped key-pecks) and of clearly distinguishing between positive and negative reinforcement (cf. chapter 5), continued study of preaversive and preappetitive stimuli showed that such an account would be an oversimplification (e.g., Kelleher, Riddle, and Cook, 1963; Azrin and Hake, 1969; Blackman, 1977).

For example, Figure 9-8 shows that shock level and baseline response rate

jointly determine whether preaversive stimuli suppress or enhance a rat's lever pressing reinforced with food (Blackman, 1968a). In the presence of red light and noise, lever presses were reinforced according to a DRL 15-sec schedule with LH 5-sec (i.e., a press was reinforced only if emitted within 15 to 20 sec of the last press); in the presence of white light and no noise, lever presses were reinforced according to FI 20-sec with LH 5-sec. In these multiple DRL FI schedules, the DRL component maintained about 5.5 responses per minute and the FI component maintained about 32 responses per minute. Later, occasional one-minute tones that preceded brief shock were added; shock level was varied to determine the relation between shock magnitude and degree of suppression. Figure 9-8 (left) shows absolute rate of responding in tone (the preaversive stimulus) as a function of shock level. In the FI component, response rate consistently decreased with increasing shock level. In the DRL component, however, response rate increased at low shock levels and decreased only at higher shock levels. Figure 9-8 (right) shows the same data converted into a suppression ratio, the change in response rate expressed relative to baseline response rate.

The different effects for FI and DRL performance indicate how behavioral phenomena may vary as a function of baseline conditions. Superimposing preaversive stimuli on reinforced responding must be treated as a respondent-conditioning procedure, but theoretical accounts of respondent conditioning must be prepared to take into account the finding that this procedure has opposite effects depending upon the baseline performance upon which it is superimposed. Similar types of interactions are critical to the analysis of the effects of drugs on behavior (psychopharmacology: e.g., Dews, 1970a); there again, the effect of a drug may vary considerably depending on the baseline performance.

Figure 9-9 illustrates how the effects of a preaversive stimulus may depend jointly on baseline response rate and baseline reinforcement rate (Blackman, 1968b). It is difficult to specify equivalent amounts of suppression with different baseline response rates (e.g., should the measure be absolute reduction, in responses per minute, or relative reduction, in percent?). Instead of examining degree of suppression the experiment examined recovery from complete suppression. Four rats served in each of nine groups that differed in baseline response rate and baseline reinforcement rate. Noise was established as a preaversive stimulus with then recorded after shock was discontinued. Baseline reinforcement rates were determined by VI scheduling: VI 30-sec (120 reinforcements per hour), VI 60-sec (60 reinforcements per hour), or VI 90-sec (40 reinforcements per hour). Baseline response rates were determined by adding to a particular VI schedule a requirement on response rate: DRL 5.0 sec (a response is reinforced only if emitted 5 or more seconds after the last response), DRP 2.0 to 4.0-sec (a response is reinforced only if emitted between 2 and 4 seconds after the last response), and DRH 0.3-sec (a response is reinforced only if emitted within less than 0.3 seconds of the last response). Strictly, these might be called tandem VI DRL, VI DRP or VI DRH schedules (cf. chapter 8). The different rate requirements generated different baseline response rates: about 10, 35 and 80 responses per minute, respectively, for the DRL, DRP and DRH requirements. Figure 9-9 shows that recovery from




Figure 9-8. Absolute rate of responding (left) and suppression ratio (right) in the presence of a preaversive stimulus as a function of the level of shock that followed the preaversive stimulus. The data shown are from one of three rats whose lever pressing was maintained by multiple DRL FI schedules of food reinforcement. With FI responding, the preaversive stimulus suppressed responding at all shock levels; with DRL responding, however, the preaversive stimulus enhanced responding at low shock levels and suppressed it only at higher shock levels. (Adapted from Blackman, 1968a, Figure 2.)

suppression in the presence of preaversive stimuli occurred more rapidly the lower the baseline response rate and the higher the baseline reinforcement rate.

The effects of preaversive stimuli are determined not only by properties of baseline performance, but also by properties of the schedule of stimulus presentation. For example, suppression varies with both the duration of a preaversive stimulus and how often the stimulus is presented within a session. The degree of suppression may depend in part on how much the reduced rate of responding affects rate of reinforcement: less suppression occurs if the reduction in responding greatly reduces the number of reinforcers in the session, and more suppression occurs if the reduction in responding only slightly affects the number of reinforcers (Stein, Sidman, and Brady, 1958; see also Meltzer and Brahlek, 1970; J. B. Smith, 1974). As with the simpler instances of respondent conditioning earlier, the effects of preaversive and preappetitive stimuli on operant behavior depend not simply on the pairings with aversive or appetitive stimuli, but rather on the differential correlation of these stimuli. The point is illustrated in Figure 9-10 (see also Figure 9-3), which shows suppression of a rat's reinforced lever pressing produced by various combinations of shock probabilities in the presence and absence of preaversive stimuli (Rescorla, 1968; see also Davis and McIntire, 1969). For example, if 40% of the preaversive stimulus presentations are paired with shock (probability of shock given CS = 0.40), a range of effects from complete suppression to no suppression at all can be obtained, depending on the probability of shock when the preaversive stimulus is absent. (The differential correlations are sometimes spoken of in terms of predictive value: the preaversive stimulus is said




Figure 9-9. Recovery from suppression in the presence of a preaversive stimulus as a joint function of baseline response rate and baseline VI schedule. With groups of four rats in each condition, three different response-rate contingencies (DRL, DRP and DRH) were added to each of the three VI schedules, so that the nine groups differed in both response rate and reinforcement rate. After complete suppression was obtained in the presence of a preaversive stimulus, recovery from preaversive stimulus. The data points show the mean trials to reach 50% recovery as a function of baseline response rate with the VI schedule as a parameter. Responding recovered more rapidly (trials to recovery were fewer) with lower than with higher baseline response rates and with shorter than with longer VI schedules (i.e., with higher than with lower reinforcement rates). (Adapted from Blackman, 1968b, Figure 1.)

to have some predictive value when the probability of shock is higher in its presence than in its absence, and no predictive value when shock is as probable in its presence as in its absence.)

The phenomena occurring when one stimulus signals another stimulus are phenomena of stimulus control: responding in the presence of the signalling stimulus is affected by its relation to the stimulus that it signals. We cannot summarize these effects in any simple way. In the classical Pavlovian cases of respondent conditioning, it appeared as if one stimulus came to function as a substitute for the other, but such accounts do not seem adequate to the effects of preaversive and preappetitive stimuli just considered. As usual, it is appropriate to recall that stimuli can have multiple functions, and it seems especially likely that such functions would have to be taken into account in situations that combine operant and respondent procedures. In addition to the cases in this chapter, we may recall that respondent phenomena were invoked in discussions of the maintenance of operant behavior (e.g., the two-factor theory of avoidance: see chapter 5 and Mowrer, 1960). If we can demonstrate that the warning stimulus in an avoidance procedure becomes aversive as avoidance responding develops (e.g., Hoffman




Figure 9-10. Suppression in the presence of a preaversive stimulus (CS) with different shock probabilities in the presence and absence of the preaversive stimulus. For example, when the probability of shock after the preaversive stimulus depended was 0.4, or $p(\text{SHOCK/CS}) = 0.4$, the suppression of the preaversive stimulus, or on the probability of shock in the absence of the preaversive stimulus, or $p(\text{SHOCK/no CS})$. Effects ranged from complete suppression when the latter probability was zero, to no suppression when the probability was equal to probability. Note that the same data are plotted in both halves of the figure: $p(\text{SHOCK/CS})$ is the parameter on the left, and $p(\text{SHOCK/no CS})$ is the parameter on the right. In this suppression ratio, baseline equals 0.5. (Adapted from Rescorla, 1968, Figure 3.)

and Fleshler, 1962), it still is one thing to say that its aversiveness comes about through respondent processes and another to say that this aversiveness is necessary to the acquisition of avoidance responding. We are perhaps still a long way from resolving such issues.

RECAPITULATION: A BEHAVIOR PARADIGM

We have considered a variety of behavioral procedures and findings: phenomena of elicitation, reinforcement, discrimination, and conditioning, among others. Before we apply these concepts to more complex types of behavior, such as language, a review seems appropriate. We now consider a paradigm in terms of which many relations in the preceding chapters can be summarized. A paradigm is a model, a graphic or symbolic representation of relations among events and concepts. Ours will represent possible relations among stimuli and responses, and therefore may provide a more formal organization for some of the phenomena considered earlier. The form of the paradigm is: $SD(R:SC)$, where SD is a *discriminative stimulus*, R is a *response*, and SC is a *contingent stimulus*. The terms within the parenthesis represent the relation between R and SC ; the parenthesis defines a *contingency*, the effect of a response on the probability of the stimulus. A contingent stimulus, therefore, is not necessarily a consequence of the response (e.g., it might be a stimulus prevented by the response, as when a response prevents shock in

avoidance). A contingent stimulus might be a stimulus delivered in a *stimulus-presentation* operation (in this case, the response has no effect on stimulus probability), or it might be a stimulus in the *consequential* operations of reinforcement or punishment (in these cases, the response raises or lowers stimulus probability).

The following development builds the paradigm from its components. Consider first the simplest experimental operation beyond mere observation, stimulus-presentation (chapter 3). We represent this operation as: (*SC*). No signal or warning stimulus precedes the contingent stimulus and no response need occur before the contingent stimulus is presented. Suppose we simply present food pellets to the rat from time to time, independently of behavior. Each food-pellet delivery is an instance of stimulus presentation, and we may watch to see what pellet deliveries do to the rat's behavior. It might eat, salivate, and perhaps become more active. (*Contingent stimulus* is justified as a name for the stimulus, because this operation is the limiting case of contingencies; just as it is useful to call zero a number, it is useful to be able to refer to this case as a special kind of contingency, even though the response has no effect on stimulus probability.)

Assume now that we change the situation so that a food pellet is delivered only if the rat rears up on its hind legs. Now the contingent stimulus *SC*, the food pellet, depends on a response *R*, the rat's rearing up on its hind legs. We represent this operation as: (*R:SC*). This relation is, of course, a *contingency*; we translate it as *the effect of response R on the probability of stimulus SC*. In this instance, food will not be delivered if the rat does not rear up on its hind legs; without this response the probability of food delivery is zero. When the response occurs, it produces the pellet, or in other words, raises the likelihood of the stimulus. This is only one example of a contingency. Responses can either raise or lower the probability of stimuli; they can turn things off as well as turn things on (chapters 4 and 5). In either case, the contingency represents a consequential operation; the contingency describes the kind of effect the response has on the environment. (The consequence of a response can also be a change in contingencies for that or another response, but this level of complexity is not critical to the present development.)

Let us now add stimulus control: we superimpose a discriminative stimulus, *SD*, on either of the other simpler operations. Suppose we turn on a light in the rat's chamber just before we deliver food. We represent this operation as: *SD(SC)*. In the presence of light, a food pellet is delivered, independently of behavior. This relation is familiar as an example of Pavlov's method of conditioning (this chapter). When Pavlov presented dogs with light and then food, the dogs began to salivate after presentation of light as well as after presentation of food.

The last example superimposed *SD* on the operation of stimulus presentation. But *SD* can also be superimposed upon the consequential operations of reinforcement or punishment (chapter 7). This time, let us turn on the light only when we will deliver a food pellet if the rat rears up on its hind legs; at other times, the light is off and no pellets are delivered. Our paradigm is now expanded to its full form: *SD(R:SC)*. In the presence of the light, *SD*, the rat's response, *R*, of rearing up on its hind legs raises the likelihood of or produces a stimulus, *SC*, the food pellet. In these circumstances, we might expect the rat to rear up on its hind legs more often when the light is on than when the light is off.

To recapitulate, the several operations may be represented as follows:

(SC)	Stimulus-presentation operation
(R:SC)	Consequential operation
$SD(SC)$	Stimulus-control operation superimposed on stimulus-presentation operation
$SD(R:SC)$	Stimulus-control operation superimposed on consequential operation.

This organization parallels the discussion of experimental operations in chapter 2; the reader may wish to review the experimental examples presented there.

In examining any experimental procedure in the analysis of behavior, it is often a useful exercise to describe the procedure (or, if it has separate stages, its various components) in terms of the appropriate operations. Sometimes a proper description requires a combination of operations. For example, assume that a rat's lever presses in light produce tone, and its lever presses in tone produce food pellets. Both stages can be represented as $SD(R:SC)$, a stimulus-control operation superimposed on a consequential operation. The response R , the lever press, is common to both stages, but different stimuli serve as SD and as SC in each stage. In the first stage, light is SD and tone is SC ; in the second stage, tone is SD and food is SC . The tone illustrates that a stimulus that serves as a contingent stimulus in one part of a procedure may serve as a discriminative stimulus in another part of the procedure (chained schedules: chapter 8).

Kinds of contingent stimuli

We have defined contingent stimuli in terms of their relations to responses, and the preceding account used the rat's food pellet as an example of a contingent stimulus. But stimuli are of various sorts, and the nature of the stimulus can determine its effects. In each of the preceding cases, we could substitute electric shock for the rat's food pellet. We could present the shock independently of the rat's behavior, (SC); we could present it only after some response, such as a lever press, ($R:SC$); we could present it only in the presence of some other stimulus, such as a warning tone, $SD(SC)$; or we could present it after some response, such as a lever press, but only if this response occurred in the presence of some other stimulus, such as a warning tone, $SD(R:SC)$. We would not expect the behavior in these cases to be the same as that with food pellets as contingent stimuli.

It is convenient, for some purposes, to be able to distinguish among different types of contingent stimuli. There are stimuli, sometimes called *appetitive*, *rewarding*, or *reinforcing*, that organisms will work to produce or to remain in the presence of: food, entertaining company, money, among many others. There are other stimuli, sometimes called *aversive*, *noxious*, or *punishing*, that organisms will work to remove or to stay away from: electric shock, dull company, extremes of heat or cold, among many others. In both cases, some stimuli have obvious biological significance while others appear to have acquired their significance during the organism's lifetime. We have seen that it is difficult to classify any stimulus unambiguously as a reinforcer or a punisher (chapters 4 and 5). And, having admitted appetitive and aversive stimuli as classes of contingent stimuli, we must

also recognize that relatively neutral or insignificant stimuli may enter into contingencies. If the rat lifts its paw and touches the chamber wall, for example, this movement is a response, and contact with the wall is its consequence. Obviously, no stimulus is likely to be completely without significance, and we must recognize that these classes represent points or regions in a continuous range of stimulus types rather than three discrete categories, and that the designation of particular stimuli can change as a result of such operations as deprivation.

Kinds of contingencies

A response may lower as well as raise the probability with which events occur. Suppose that shock is delivered to a rat at the end of every 10-second period unless the rat presses a lever. The probability of shock is 1.0 at the end of 10 seconds if the rat has not pressed the lever, but it is zero if the rat has pressed the lever; the lever press in any 10-second period reduces shock probability from 1.0 to zero. Because a contingency is defined as *the effect of a response on the probability of a stimulus*, we may distinguish among contingencies in terms of whether the effect is an increase or a decrease in probability.

Consider some examples. In the last case, in which the rat's lever presses avoid shock, the stimulus is aversive, and the response reduces its probability. If, however, the rat's lever presses produce shock, the contingent stimulus is again aversive, but the response this time raises its probability. Similarly, a rat's lever presses may either raise or lower the probability that an appetitive stimulus such as a food pellet will be delivered.

Changes in probability are, of course, not limited to all-or-none cases, +1.0 and -1.0. For example, a rat's lever press may produce a food pellet only one-tenth of the time (e.g., ratio schedules: chapter 8). In this case, the response raises stimulus probability only from 0 to 0.1. Similarly, a rat's lever press may prevent only half of the shocks arranged at regular time intervals. In this instance, the probability of shock decreases from 1.0 to only 0.5. In other words, contingencies can be represented continuously on a scale ranging from -1.0 through 0 to +1.0. (We do not know enough about contingencies to say whether a given change in probability has the same effect no matter what the starting probability. In fact, we might suspect that a change in stimulus probability from 0 to 0.5 would not be equivalent to one from 0.5 to 1.0, even though both are increases of 0.5.)

It is important to note that the order of terms in the symbolic representation of a contingency does not necessarily represent the order in which stimuli and responses occur in time. For example, we may write ($R:SC$) for the case in which a rat's lever presses turn off or prevent electric shocks, but this means that shocks are delivered *only* in the absence of responses; if a response occurs, it is *not* followed by shock. A contingency represents a *relation* between responses and stimuli; we are interested in the succession of responses and stimuli in time, but must not confuse this temporal succession with the order in which symbols appear on the printed page. We write ($R:SC$) for the case in which a rat's lever press produces food pellets not so much because the lever press will be followed by a food pellet, but because it is important to represent the environmental consequences of lever pressing.

In classifying contingent stimuli, we recognized relatively neutral or insignificant stimuli as well as appetitive or aversive stimuli. The same consideration is appropriate in the case of contingencies. The limiting case of a contingency is one in which a response has no effect on stimulus probability. For example, we may deliver food or shock to a rat independent of its lever presses. These are simple stimulus-presentation operations, but the notation ($R:SC$) emphasizes the relation between stimulus-presentation operations and consequential operations. Thus, ($R:SC$) could represent the response-independent delivery of appetitive stimuli such as food pellets, or aversive stimuli such as electric shocks, or relatively neutral or insignificant stimuli such as lights or tones.

Classification of contingencies and contingent stimuli

But a limitation to the paradigm must be recognized. In some contingencies, it is difficult to decide whether the change involves presentation or removal of a stimulus. For example, is the response-produced change from a cold to a comfortable environment to be counted as positive because heat has been added to the environment, or as negative because the sensory effects of cold have been terminated? Contingencies involve *changes* in the environment, and we must recognize that there will be ambiguous cases. We might even go further, and consider whether the distinction between positive and negative reinforcement or between positive and negative punishment is worth maintaining. The important distinction is that reinforcement increases responding whereas punishment decreases it; whether these increases or decreases come about through presenting stimuli or taking them away is perhaps less critical (cf. chapter 5).

Kinds of contingencies and contingent stimuli are summarized in Table 9-1, in which they are shown combined with a discriminative stimulus, SD . These instances are shown with some representative names applied to them in the psychology of learning. The specific procedures are not exhaustive, and the list of names is therefore incomplete. For any procedure in the psychology of learning, it is instructive to locate it or its various stages among the classes in Table 9-1.

Relations among responses

An important property of a contingent stimulus is the behavior that it occasions. We must distinguish the responses involved in contingencies from responses occasioned by the contingent stimulus. For example, if the contingencies involved production of food pellets by a rat's lever presses, the lever press is specified in the contingency; other responses such as taking the food in the mouth, chewing, salivating, and swallowing are occasioned or elicited by the pellet. The effects of contingencies may depend in part on the relation between these response classes.

First, we may examine the relative probabilities of these responses in situations in which both are possible. If we were interested in the responses of running and of eating, we might see which the rat engaged in when both a running wheel and food pellets were freely available. We may expect an opportunity to run to reinforce eating if running is more probable than eating, but not otherwise.

TABLE 9-1

Kinds of contingencies and contingent stimuli. Entries are representative classes of experimental procedures; those in parentheses illustrate cases in which a discriminative stimulus is superimposed on the contingency.

Type of contingent stimulus	Type of contingency		
	<i>Response raises probability of stimulus</i>	<i>Response does not affect probability of stimulus</i>	<i>Response lowers probability of stimulus</i>
<i>Appetitive, rewarding, or reinforcing</i>	Positive reinforcement (operant discrimination)	Stimulus presentation (respondent conditioning)	Negative punishment (omission training)
<i>Relatively neutral, or insignificant</i>	Sensory consequences (latent learning)	Stimulus presentation (sensory preconditioning)	Sensory consequences (latent learning)
<i>Aversive, noxious, or punishing</i>	Positive punishment (discriminated punishment)	Stimulus presentation (defensive conditioning)	Negative reinforcement (discriminated avoidance)

Conversely, we may expect enforced running to punish eating if running is less probable than eating, but not otherwise (chapter 4 and Premack, 1959). In other words, the effect of a contingency may depend on the relation between these response classes. In this context, we may formulate the study of motivation as the study of the conditions that determine the effectiveness of contingent stimuli as reinforcers or as punishers.

The responses in a contingency and those occasioned by contingent stimuli may differ in similarity or compatibility as well as in relative probability. For example, if a pigeon's pecks on a key produce food, pecking on a key and pecking at food are similar in topography though they differ in location. But if the pigeon's pecks avoid shock, wing flaps or other responses produced by shock may have little in common with pecking. *In general, reinforcement is likely to be more effective if the reinforced response is similar to or compatible with responses occasioned by the reinforcer than otherwise, whereas punishment is likely to be more effective if the punished response is dissimilar to or incompatible with responses occasioned by the punisher than otherwise.* In the case of reinforcement, consider avoidance as an example: if a rat's response can avoid electric shock, a response like jumping that has much in common with shock-elicited responding may be easier to establish than a response like lever pressing that is incompatible with shock-elicited jumping. In the case of punishment, shock as an aversive stimulus, which may elicit vocalization and biting in the rat, is likely to be more effective in punishing lever pressing, a response somewhat independent of vocalization and biting, than in punishing either vocalization or biting themselves, responses elicited by the very stimulus used as the punisher (chapter 5).

Conclusion

We expressed experimental operations in terms of relations among stimuli and responses. Except for the mere observation of behavior, the presentation of stimuli

is the simplest operation, (SC). But, because responses can have environmental effects, we must also examine consequential operations, which are specified in terms of contingencies, ($R:SC$). A contingency is the effect of a response on the probability of a stimulus. In addition, either stimulus-presentation operations or consequential operations can be signalled, i.e., can take place in the presence of particular stimuli. Thus, we must also be concerned with stimulus-control operations, whether these are superimposed on stimulus-presentation operations, $SD(SC)$, or on consequential operations, $SD(R:SC)$. These various operations can involve contingent stimuli that are appetitive, aversive, or relatively insignificant. Further, the contingencies into which they enter can involve increases in stimulus probability, decreases in stimulus probability, or no change in stimulus probability. The last special case of contingency is equivalent to a stimulus-presentation operation: ($R:SC$) = (SC).

These classifications do not guarantee that any stimulus or any response will have only a single function; a stimulus in a contingent relation with one response may be in a discriminative relation with another, and a response elicited by one stimulus may be in contingencies with other stimuli (for example, the mother's presence may be a contingent stimulus, when she comes at the infant's cry, and a discriminative stimulus, when the infant learns that things happen in the mother's presence that do not happen in her absence; the infant's cry may sometimes be elicited by events, such as painful stimuli, while at other times it may occur because such events as the mother's presence are its consequence). We come to understand behavioral situations by separating the various stimulus and response relations that enter into them. That is the business of an experimental analysis. In studying behavior, we must analyze situations to determine which features of an operation were important and which aspects of responding were affected by it.

On close examination, some distinctions implied by the paradigm seem to diminish in importance. In the analysis of behavior, classifications of events often have fuzzy boundaries, and distinctions may then become arbitrary. We noted such a case with respect to presenting or removing stimuli (e.g., is food effective as a reinforcer by virtue of its presentation or because it terminates stomach contractions or other events that may be correlated with hunger?). In the final analysis, we may completely discard the distinction between presenting and removing stimuli, recognizing that all consequential effects of responding can be characterized simply as environmental changes. We might then note that every procedure takes place in some environment, and so we might next get rid of the discriminative stimuli, observing as we did so that SD can be incorporated into the definition of the response. Thus, if a pigeon's key pecks produce food in the presence of green but not red light, the response of the contingency can be defined to include only those pecks emitted in green, and we can simplify the paradigm to ($R:SC$), where R represents key pecks in the presence of green and SC represents food (chapters 6 and 7). If later our analysis shows us that we cannot even define response classes independently of contingencies, we might ultimately decide that we can completely dispense with the paradigm. But it has been useful along the way to make some distinctions explicit. We can, so to speak, throw away the ladder after having climbed it.

Verbal Behavior

FUNCTIONAL PROPERTIES OF VERBAL BEHAVIOR

Formal classes of verbal behavior

Echoic behavior

Transcription behavior

Textual behavior

Dictation behavior

Relations among the formal classes

Tacts: The relation between language and environment

Abstraction

The extended tact

Tacts of private events

Intraverbal behavior

The consequences of verbal behavior

The mand as a class of verbal responses

The role of instructions

The listener's or reader's behavior

Autoclitic processes: Verbal behavior conditional upon verbal behavior

Descriptive autoclitics

Qualifying and quantifying autoclitics

Relational autoclitics

Interaction between verbal behavior and environment

Our language words have diverse sources. *Verbal*, through Latin, and *word*, through Old English, are derived from a common Indo-European root, *wer-*, to speak. The Germanic *spek-* or *sprek-*, from which comes the German *die Sprache*, speech or language, leads to the English *speak* and *speech*. The Greek *legein*, to speak, and *logos*, word, lead to *lexical*, *legible*, and such relatives as *logic* and *intelligent*. Latin provides *language* and *linguistics*, through *lingua*, tongue, and *vocal* and *vocabulary*, through *vox*, voice.

Language is behavior. But our everyday vocabulary of language includes many assumptions that stand in the way of a behavioral account. For example, consider the common term *word*. When we speak of words in language, we seldom bother to distinguish spoken words from written words. Yet speaking a word is not the same as writing it, and the two responses are likely to occur in different circumstances. Even worse, we often speak of *using* words, as if these were manipulable objects instead of behavioral events.

Particular words are uttered or written in particular circumstances. We do not ordinarily say that we utter nouns in the presence of relevant objects, or that we utter sentences in the presence of relevant events. Instead, we tend to say that words refer to objects, or that sentences are about events. There is good reason for this usage in the everyday vocabulary; as we shall see, it is appropriate to the effect of language on a listener or a reader. But this vocabulary may be misleading when applied to an analysis of a speaker's or a writer's behavior.

An analogy may clarify the point (cf. MacCorquodale, 1970). Some ancient accounts of vision assumed that vision depended on emanations from the eye making contact with the object, rather than on light from an object entering the eye. The direction of effect, from eye to object or from object to eye, was critical

to subsequent progress in vision and optics. The distinction between the two directions is sometimes implied in the vocabulary of vision (e.g., as when we distinguish between something seen and something looked at), but at least the vocabulary does not prejudice an account of visual phenomena in one direction or the other.

This is not the case in the vocabulary of language. We speak of language as if it were directed toward environmental objects or events: we say that words or sentences refer to, deal with, speak of, call attention to, or are about things. The language of reference implicitly includes the direction from verbal behavior to environment. A vocabulary emphasizing the opposite direction is not obviously available in the everyday language of verbal behavior. We must therefore consider the possibility that the everyday vocabulary of language has prejudiced the ways in which we can analyze verbal behavior.

Another part of the everyday vocabulary that complicates the task of analyzing verbal behavior is the language of meaning. We say that utterances convey meaning. But dictionaries do not contain meanings of defined words; they just contain other words. It is not useful to say that an utterance conveys a meaning unless something also can be said about the sources of the meaning. The trouble is that meanings are not identifiable as events. In fact, the concept is derived from verbal behavior, and it is therefore inappropriate to use it in turn to explain verbal behavior.

Accounts of language in terms of such concepts as meaning are, inevitably, theories of language. To be effective, theories must classify the events with which they deal. A primary task of the analysis of language, therefore, is classifying verbal behavior. But if meaning and related concepts provide an ambiguous basis for classifying verbal behavior, what properties of verbal behavior should take their place?

Grammatical classifications exist, but the grammatical properties of the words in a sentence will not help us to distinguish among the circumstances in which a sentence might have been generated or the consequences its production might have for a speaker, a writer, a listener, or a reader. For example, a sentence such as "The psychologist talked to the chimpanzee" has the same grammatical properties whether uttered by the observer of a human-simian conversation, written as memorized by a student taking an examination, overheard by a passing clinician in a psychiatric ward, or simply encountered as an arbitrary example in a text. In all of these cases, we would class *psychologist* and *chimpanzee* as nouns and would note the different relation each has to the verb *talked*; yet the cases are worth distinguishing, even though they are grammatically identical.

Describing the preceding sentence in terms of grammatical properties of its components deals with the structure of verbal behavior. But distinctions made on the basis of the circumstances in which a sentence can occur are distinctions of function. A functional account of verbal behavior is based on how verbal responses work: conditions under which they occur and consequences they may have. As with nonverbal behavior, we will not have to choose between structural and functional accounts of language, because the two accounts complement each other. Unfortunately, verbal behavior has been controversial in the history of psychology, and structural and functional accounts have often been pitted against each other as

if they were incompatible instead of complementary (cf. Titchener, 1898; Catania, 1973c; Segal, 1977). This text attempts to deal consistently with both kinds of approaches. The present chapter outlines a functional account of verbal behavior (Skinner, 1957).

FUNCTIONAL PROPERTIES OF VERBAL BEHAVIOR

Verbal responses can be distinguished by the occasions on which they occur and the consequences they produce. They can be occasioned by either verbal or nonverbal stimuli, and they can have either verbal or nonverbal consequences. For example, a child might say "apple" in the presence of either the written word or an actual apple; as a consequence of saying "apple," the child might be given either the verbal reply, "right," or the apple. The following account argues for the importance of these kinds of distinctions. First, we consider some relatively simple formal relations, in which verbal behavior is reproduced or imitated in either the vocal or the written mode (e.g., as in echoing what someone has said); note the distinction between the general term, *verbal*, and the term specific to spoken language, *vocal*. These cases will be followed by a treatment of verbal behavior under the control of nonverbal stimuli (e.g., as in naming different objects in the environment) and under the control of verbal stimuli (e.g., as in giving definitions of terms). The account will also deal with audiences as occasions of verbal responses, the behavior of listeners, and some consequences of verbal behavior that may contribute to its maintenance. Finally, an exploration of more complex verbal processes (e.g., as in assertion and negation) will set the stage for discussing some structural properties of verbal behavior.

Formal classes of verbal behavior

Formal verbal relations are defined by the correspondence between verbal stimuli and verbal responses (in the colloquial vocabulary, we would say that the words; whether stimuli or responses, are the same). We will restrict our attention to vocal and written verbal behavior, although we could easily extend the account to other modalities (e.g., Braille or sign language). One of the most elementary verbal functions is the reproduction of verbal behavior: we repeat what others have said, or copy what others have written. Thus, the cases include the various combinations of vocal or written stimuli and vocal or written responses.

Echoic behavior. Imitation of the acoustic properties of vocal stimuli appears relatively early in human infants' acquisition of speech. This class of verbal relations is called *echoic*. When a parent says "mama" and the child repeats "mama," the child's response is echoic to the extent that it is occasioned by the parent's utterance. Even though the stimulus and the response have common properties, this verbal relation is not simple. The stimulus is a complex sound pattern. The response consists of coordinated movements of the vocal apparatus. These movements produce a sound pattern similar, in important acoustic properties, to

the stimulus, but these movements are not themselves sounds. The properties of the heard sound pattern and the movements of the vocal apparatus that echo that sound pattern are only indirectly related. How does the child who hears sounds like "mama" or "dada" come to make precisely those vocal responses heard as the same sounds by the parents?

The speaker's articulation of sounds, the physical properties of the acoustic stimulus, and the listener's response to these sounds are intricately related. For example, many consonants of English (*p, b, t, d*) cannot even be produced unless they are accompanied by a vowel, and their physical properties vary as a function of the context of sounds within which they occur (e.g., Lane, 1965; Liberman, 1970). The capacity for echoic behavior must depend in some way on the shaping of vocal apparatus movements by the acoustic consequences of vocalization. To the extent that these consequences engender and maintain vocalization, we may say that these consequences are reinforcing, but we must recall that the concept of reinforcement describes and does not explain. Thus, we may regard native-language speech sounds as reinforcing relative to nonnative-language speech sounds, because native-language sounds are ordinarily maintained in an infant's spontaneous vocalizations whereas nonnative-language sounds gradually disappear. Native-language speech sounds are presumably discriminated and perhaps acquire their reinforcing properties because they often accompany the activities of important caregivers in the infant's environment (Eimas, Siqueland, Jusczyk, and Vigorito, 1971; Trehub and Chang, 1977).

Echoic behavior does not constitute the major part of human verbal behavior.¹ Nevertheless, echoic behavior does not simply accompany the acquisition of human language and then vanish; rather, it persists in the behavior of mature speakers. For example, we may repeat a number given to us by a telephone operator, or the name of someone who has just been introduced to us. The acoustic correspondence of the vocal stimulus and the vocal response characterizes this class of verbal behavior.

The unit of echoic behavior may vary in size from individual speech sounds, or *phonemes*, to extended phrases or sentences. For example, the echoic production of individual speech sounds may be important in initially learning to pronounce words in an unfamiliar language, or in generating rhymes or alliteration in poetry (B. H. Smith, 1968; Skinner, 1972). On the other hand, the echoic production of extended phrases or sentences may be important in dramatics, as when an actor repeats the lines whispered by a prompter, or on ritual occasions, as when a bride and groom repeat the phrases of a marriage vow spoken by a minister. The unit of verbal responding is not defined by its size; rather, the unit is defined by the functional relation between stimuli and responses, and is demonstrated when the verbal response can be shown to occur in verbal behavior independently of other verbal units.

Transcription behavior. Verbal stimuli can also correspond to verbal responses when both the stimuli and the responses are written. In such cases, the

¹Echoic responding occurs naturally among some nonhuman organisms, such as parrots and myna birds (cf. Mowrer, 1950), but the units of echoic behavior in animals are presumably different from those of human speakers.

behavior is called *transcription*. For example, we may copy a number from the telephone book, or copy an author and title in preparing a bibliography. But, as in echoic behavior, the movements in producing a word may not be related in any simple way to the properties of the word as a stimulus; writing a word is different from looking at a word. The correspondence need not even be in physical properties of the stimuli and responses. In the form of individual letters, a type-written text may have little in common with a handwritten copy. Nevertheless, we can ordinarily judge whether a transcription is accurate in spelling, word order or punctuation.

Just as the unit of echoic behavior may vary from individual phonemes to entire phrases or sentences, the unit of transcription also may vary from individual characters to extended passages, depending on the circumstances in which the behavior occurs. A child learns to copy individual letters before learning to copy entire words. In doing so, the child learns the correspondences between arbitrary visual forms, such as the printed and script *a* in upper and lower case. We may be unable to provide unique physical specifications common to the letter *a* in its various forms and to no other letters, but the problem is not different in kind from that of specifying the unique acoustic properties of a given phoneme as it occurs in different contexts (cf. Gibson, 1965).

It is important to distinguish transcription from copying in the pictorial sense. A skilled Asian calligrapher may be able to produce an accurate and readable copy of an alphabetic text even though unfamiliar with the European language in which the text is written, but it would be inappropriate to speak of such copying as verbal. The distinction is based on the behavioral units in the two kinds of copying. The critical features of the calligrapher's copying are the geometrical properties of the letters in the text and the marks produced by the calligrapher's strokes. The critical features of transcription, however, are the verbal units, such as letters, words and phrases, in the original text and its copy.

Except for their respective vocal and written modes, echoic behavior and transcription behavior are formally similar. It is tempting to assume, perhaps because of the relative ease with which it is learned, that echoic behavior is simpler than transcription. Children ordinarily acquire echoic behavior early, even without specific instruction, but they do not ordinarily acquire transcription unless it is explicitly taught. Acquiring transcription behavior may also involve other kinds of responses, as when a child is taught to say a word aloud while copying it.

Pure transcription, in the sense of transcription unaccompanied by other responses to the text (such as vocal or subvocal responses), probably occurs only rarely. A skilled typist, for example, may sometimes transcribe a text while not responding verbally to the text in other ways (as when listening to a conversation elsewhere in the office); in such circumstances, the typist is unable to report the text content even though it was accurately transcribed. A particular stimulus may set the occasion for a variety of responses, and a particular response may be occasioned by a variety of stimuli. The integrity of this functional classification is not diminished by the possibility that transcription is sometimes accompanied by other kinds of verbal behavior.

Textual behavior. When a written verbal stimulus sets the occasion for a corresponding vocal verbal response, the behavior is *textual*. Thus, we may say aloud a number we have found in a telephone book, or read a bedtime story to a child. In textual behavior, the arbitrary correspondence between verbal stimuli and verbal responses is more obvious than in either echoic or transcription behavior, because the stimuli and the responses are in different modes. As with transcription, this behavior is usually taught explicitly, and some controversies in such teaching are based on assumptions about the behavioral unit appropriate to various stages of instruction (e.g., whether, in teaching reading, the teacher should begin with individual letters, syllables, or entire words: cf. Gleitman and Rozin, 1973).

As with the other formal classes, textual behavior must be distinguished from other kinds of responses to written verbal stimuli. For example, if a sign says "Stop," there is a difference between reading the word aloud and stopping; only the former is textual. In the mature reader, textual responses become less important than other kinds of responses to written verbal stimuli. Vocal responses diminish in magnitude, become subvocal, and perhaps disappear completely as a child becomes a proficient reader. Reading is behavior, but textual responses, as defined here, are not equivalent to reading. For example, a parent, while reading a bedtime story to a child, may discover that the most recent page, read without attention, cannot be remembered; yet the child's reaction may indicate that the page had been read without any serious errors or omissions. This example indicates some problems in the colloquial vocabulary, which does not distinguish between reading that is simply the saying of the words on a page and the kind of behavior we call reading for understanding. Some advantages of a behavioral analysis come from making such distinctions. Textual behavior is defined in terms of vocal responses to visual verbal stimuli; these vocal responses may eventually diminish, as when a child's lips move in silent reading, but they continue to be textual responses. Because reading for understanding must include other behavior along with or instead of vocal or subvocal speech, it is more than simply textual behavior.

Dictation behavior. Just as a written verbal stimulus can set the occasion for a vocal verbal response, a vocal verbal stimulus can set the occasion for a written verbal response. This class of verbal behavior is called *dictation* (here we emphasize the behavior of the listener, who takes the dictation, rather than that of the speaker, who dictates). For example, we may write a number spoken to us by a telephone operator, or take notes at a lecture. The units of dictation are typically entire words or phrases, but individual letters may also serve as units (e.g., as when children are taught the written alphabet, or as when an unusual name is spelled out for a stenographer).

As with textual behavior, dictation involves stimuli and responses in two different modes. Some of its unique properties follow simply from the relatively permanent record, in the written text, produced by dictation. In addition, occasions for dictation are limited relative to textual behavior, because writing implements, unlike the vocal apparatus, are not parts of the human anatomy. Perhaps for this reason, there is little temptation to pursue the possible covert manifestations of dictation behavior; we are less likely to speak of submanual writing or typing than of subvocal reading (the motor theory of consciousness had

argued that thought was merely language behavior, like speech, that had been reduced in magnitude: cf. Max, 1934). Nevertheless, textual behavior and dictation behavior are formally similar, and either can be accompanied by other kinds of verbal responses occasioned by verbal stimuli.

Relations among the formal classes. This account of formal verbal classes has been limited to vocal and written stimuli and responses. It could have been extended to other language modes (such as sending and receiving Morse code, or writing and reading Braille). It is not critical, however, to provide an exhaustive list of all possible formal verbal relations. The present classes are important because they emphasize distinctions in terms of verbal stimuli and verbal responses, and these distinctions are relevant to the acquisition of language. We speak of letters and words without regard to whether they are written or spoken, but letters and words function differently as stimuli and as responses in these modes. For example, a skilled teacher who had taught a child only to name letters written on a blackboard would not then expect the child, without additional instruction, to be able to write the letters as they were spoken. The stimuli in the former task are the responses in the latter, and vice versa. Both of these relations, and others as well, are part of the child's "knowing the alphabet." But the distinctions among them cannot easily be made in the everyday vocabulary.

Perhaps these distinctions are obscured in the everyday vocabulary simply because language instruction is designed to eliminate them; to teach reading and writing is to teach equivalences between the acoustic and visual modes of language. The relations are summarized in Figure 10-1. Each mode, vocal and written, has special characteristics. For example, stress, rhythm, and intensity vary more freely in spoken than in written verbal behavior, but spoken verbal behavior is more transient than written. Nevertheless, some important properties of verbal behavior are independent of mode. For example, for the immediate consequence of being able to dial a telephone number, it may make little difference whether one hears the number as spoken by a telephone operator or reads it as written in a telephone book, or whether one repeats the number vocally or copies it down. The different consequences that do exist (e.g., the time it takes to get the number, and the likelihood of retrieving it on another occasion) may affect the caller's future behavior in obtaining telephone numbers, but such differences are not essentially verbal. Thus, distinguishing among formal classes may be relatively unimportant in analyzing the verbal behavior of a mature speaker, but may be critical in studying language acquisition, when the formal classes are not yet well established, or in interpreting language pathologies, when the deficits may be characterized by the particular verbal classes affected (e.g., the language deficits called aphasia: cf. Sidman, 1971).

Thus far, we have emphasized verbal classes defined in terms of stimulus and response modes, but other formal relations among verbal classes may be differentiated even within a single mode. For example, if a student has learned only to translate from English to German, the student may have difficulty when asked to translate in the other direction, from German to English. Learning German responses to English stimulus words is different from learning English responses to German stimulus words. (The problems are compounded when both spoken and




Figure 10-1. Relations among the four formal classes, in which S = stimulus and R = response. Similar relations can be established for any other pair of verbal modalities (e.g., spoken verbal behavior and the gestural verbal behavior of American Sign Language).

written languages are involved. Language instruction usually recognizes these distinctions; a course in conversational French is expected to emphasize the vocal mode, whereas one in scientific Russian is expected to emphasize the written mode.)

The formal classes illustrate the importance of distinguishing between verbal stimuli and verbal responses. All the examples have been variations on the same theme. Distinguishing between stimuli and responses presents little difficulty in the analysis of nonverbal behavior. If a rat presses a lever only when a light is present, we call the light a stimulus and the lever press a response; we are not tempted to reverse the terms. But in verbal behavior, the speaker's response is the listener's stimulus and the writer's response is the reader's stimulus, and confusion comes easily. Thus, we must characterize the stimulus and response components of various verbal relations with care.

Tacts: The relation between language and the environment

A *tact* is a verbal response occasioned by a nonverbal stimulus. For example, if a child learns to say "apple" in the presence of an apple, the child is said to be tacting the apple. The tact is an instance of stimulus control; a particular stimulus sets the occasion for a particular response. Through this relation, verbal behavior makes *contact* with the nonverbal environment. Tacting has much in common with naming, but it is distinguished from naming by the presence of the tacted stimulus. We may speak of naming an object that is absent, but we could not then say that the object had been tacted.

An indeterminate variety of tacts is available to the mature speaker. We tact objects (chairs and tables, pencils and books), living things (flowers and trees, birds and insects), weather conditions (rain and snow, sun and clouds), activities (walking and running, working and playing), and innumerable other features of the environment. Some tacts are general (e.g., *man*, *woman*), whereas others are restricted to relatively narrow circumstances (e.g., the name of a particular individual). The wealth of available tacts may be taken as a remarkable feature of the human capacity for language, but this wealth should not obscure the relations that define any instance of tacting. It is useful to consider how tact behavior might be established in nonhuman organisms; such examples force us to be explicit about the components that must enter into such behavior.

Let us imagine a food-deprived pigeon in a chamber with a window on which different colors can be projected. Beneath the window are three keys. When the window is lit red, pecks on the left key produce food; when the window is lit blue, pecks on the middle key produce food; and when the window is lit green, pecks on the right key produce food. In such circumstances, the pigeon will eventually peck the left key during red, the middle key during blue, and the right key during green. Is there any reason why the pigeon's performance should not be called the tacting of red, blue, and green?

It first might be argued that the stimuli are highly specific; the pigeon would not be expected to respond to red apples and red shirts and red sunsets as it responds to the red light in the window. We have already noted that the generality or specificity of the controlling stimuli does not define the tact relation. If we are in doubt about the generality of the pigeon's response to red, we may simply call the pigeon's peck on the left key a tact of a red-lit window in this particular chamber; in effect, we may treat the left-key peck as the pigeon's name for this stimulus.

What of the consequences of the pigeon's pecks? If the pigeon were not food-deprived or if food were not delivered as a consequence, the pigeon would stop pecking. Even human tacting is not ordinarily without consequences. We do not move about our environments naming everything we see, but we tact when tacting has consequences. The consequences do not define the tact, but tacting remains tacting whether maintained by approval, pay, an examination grade, or simply because it is helpful to the individual to whom we are speaking.

Finally, we may object that the pigeon has no audience. Although the pigeon's pecks may be appropriately occasioned by red and blue and green lights, the pigeon is not speaking to anyone. This difficulty could easily be rectified, by arranging signs, "red," "blue" and "green," that respectively lit up after the pigeon's pecks on the left, middle or right key. If we knew our pigeon's performance was accurate and we could not look into the chamber, we could simply watch the signs and let the pigeon tell us which light was on. The pigeon would be speaking in the language of pecks, but its verbal responses would nonetheless be tacts.

If we utter a word in the presence of an object, we do not reject this as naming if no one has been listening. An audience may be important in establishing or maintaining tacts, but it does not define the tact relation. The tact is defined by the relation between stimuli and the verbal responses occasioned by them. The likelihood of tacting may be modified by audience variables and consequences, but these variables are not criteria for whether a tact has occurred. Although the pigeon's pecks in the presence of red and blue and green may be relatively trivial instances of verbal behavior, they are legitimate illustrations of the tact relation.

Abstraction. We have emphasized the control of tacts by relatively simple environmental events. Strictly, we should say that tacts are controlled by properties of the environment rather than by particular stimuli or classes of stimuli. The property of color, for example, is the critical determinant of the verbal response, "red," whether that response is occasioned by red fire engines, red traffic lights, or red cheeks. This verbal discrimination based on a single property of a stimulus is

called an *abstraction*. The property is defined by the practices of the verbal community; it does not depend on our capacity to provide some independent physical measure. For example, no range or distribution of wavelengths can be established such that all visual stimuli within the range are called red whereas all outside the range are not. Our ability independently to specify some measurable physical dimension of the controlling stimulus is not essential to identifying a particular verbal relation as a tact.

Experiments on concept formation may be regarded as demonstrating the acquisition of tacts (cf. chapter 7 on concepts). An example is an experiment (Hull, 1920) in which learners had to master names for each member of several sets of Chinese characters (three sets are shown in Figure 10-2). A given name was consistently related to a particular radical that appeared in one character in each set, but the position and size of the radical and the configuration within which it appeared varied from one set to another. Each set of twelve characters was presented until the learner could give the appropriate name for each, and then a new set of twelve was presented. By the fifth set, learners were able to name more than half the characters upon seeing them for the first time, and sometimes they were able to give the name even though they could not sketch the radical or describe the basis of their naming.

Hull's experiment may be regarded as concerned with the acquisition of tacting. Tacting in that situation differed from tacting in natural languages in that the basis for many of our tacts cannot be so explicitly defined. For example, we cannot say exactly what properties make an object a chair. A chair may have four legs or stand on a single pedestal, it may have a flat or a contoured seat or back, and it may be constructed from many different materials. We cannot even appeal to its function, because we call some objects that cannot be sat upon chairs (e.g., a toy chair in a set of doll-house furniture).

	Radical	Characters Lists
yer	扌	殳 殳 殳 殳 殳 殳 殳
ta	弓	弧 弧 弧 弧 弧 弧 弧
ling	宀	宀 宀 宀 宀 宀 宀 宀

Figure 10-2. Radicals and their appearance as parts of Chinese characters in six different lists. Learners had to give the radical name to the character that contained it, but characters were not repeated in successive lists. Learners became able to name characters appropriately upon seeing them for the first time in new lists, and sometimes were able to do so even when they were unable to define the common radical by sketching it. (From Hull, 1920, Figure 1.)

Extremely subtle properties can control tacts. These properties may include relations among stimuli. Such terms as *above* and *below*, *near* and *far*, or *larger* and *smaller* tact properties of stimuli in relation to each other or to the speaker. Relational tacting occurs when we say that two objects are alike or are different, or when we note that one item in a set is an odd item. Such terms rarely stand alone, and we will consider their joint dependence on the relational properties of events and on other verbal responses when we later consider the verbal relations called *autoclitic*. We also may tact complex events extended in space or time, as when we discriminate among different authors on the basis of their writing styles, or as when we identify a musical piece as one by Debussy or a painting as one by Monet. At another level of complexity, we might say that Debussy and Monet have something in common, even though it would be difficult to specify common dimensions of Debussy's music and Monet's paintings. Yet to call them both impressionists may be regarded as a tact of common properties.

Sometimes the properties controlling a tact can be identified more with the speaker's behavior than with any particular feature of the environment. For example, if a painting or a musical composition or a person occasions the tact, "beautiful," this tact presumably depends on the common responses generated in the speaker rather than on any particular physical properties common to these stimuli. An interesting verbal case is the "tip-of-the-tongue" phenomenon. When we say a word is on the tip of our tongue, we are tacting the near-threshold availability of an appropriate verbal response in our own verbal behavior. Sometimes we can even report properties of the unrecalled word, such as its length or part of its spelling (R. Brown and McNeill, 1966).

The vocabulary of emotion is similarly based on complex relations among stimuli, situations, and the speaker. Tacts of love, hate, joy and sorrow, whether in oneself or in others, are based not simply on overt manifestations such as laughter or tears, but also on the various circumstances that generate the observed behavior. If this were not the case, it would be hard to imagine how a verbal community could maintain any consistency in this vocabulary; its variability is itself evidence of the subtlety of the relations tacted.

Any event or situation obviously may have many properties that might be tacted. Whether any properties are tacted and which properties are tacted will depend on other variables that may act upon the speaker. For example, we may tact the color of an apple in one circumstance and its smell in another. The situation presents no difficulties; verbal responses may be determined in multiple ways. An obvious case is when the tact is also a response to a question (e.g., "What color is this?" "What does it smell like?").

Consider again the analogous case for the pigeon. We might alter the lights behind the windows in the chamber so that the colors appeared bright, moderate, or dim. In the presence of a tone, we could then arrange that regardless of color left-key pecks produced food in the presence of bright light, middle-key pecks produced food in the presence of moderate light, and right-key pecks produced food in the presence of dim light. If the pigeon's pecks remain appropriate to color in the absence of tone, but become appropriate to intensity in the presence of tone, we might say that the pigeon was tacting color in the absence of tone and intensity

in the presence of tone. In the vocabulary of behavior, this is a conditional discrimination. The analogous case in human verbal behavior is when a question serves as a conditional stimulus (corresponding to tone in the preceding example) that sets the occasion for tacting a particular property of the many properties of a stimulus (e.g., the color or smell of an apple). The audience, previous verbal behavior, and a variety of other factors may affect tacting. In other words, tacting, like other kinds of behavior, is a class of responses that may come under the control of environmental conditions. We do not tact indiscriminately. We tact some things in some circumstances and other things in other circumstances; we also learn that in some circumstances (e.g., remarking on someone's bad breath or dandruff), it may not be tactful to tact at all.

We can also tact the temporal dimensions of stimuli, as when we say that something lasted a long or a short time. Often we respond verbally to stimuli that are no longer present, but we may best regard tacts as responses occurring in the presence of or very shortly after the environmental events occasioning them; responses occurring long after such events require special treatment because they typically include other behavior besides tacting (see chapter 13 on memory). Nevertheless, environmental events have important temporal dimensions, and these temporal dimensions can be discriminated. Sentences in different tenses may be regarded as classes of tacts controlled by temporal properties of the environment. We may say, "It is raining," "It is beginning to rain," or "It just stopped raining." Each sentence is a different tact; all are controlled by the event, rain, but they are distinguished by the different temporal properties of that event. A past-tense verbal response such as "It rained" may also be uttered under the control of other kinds of stimuli, when it is not appropriate to call it a tact. For example, it might be uttered as an echoic or a textual response, or as a tact under the control of present stimuli such as wet streets. In such a case, the response may depend on more complex verbal behavior, as when derived from, "The streets are wet; therefore it must have rained."

The extended tact. The tact is a flexible relation. In some verbal communities, the stimulus properties controlling a tact may be sharply defined. A student in a science laboratory, for example, is taught to be consistent in tacting particular apparatuses and procedures. This precision is less common in everyday discourse. We often tact properties of the behavior of our acquaintances, but the conditions under which someone is said to be warm or reserved, energetic or lazy, interesting or dull, and so on, vary considerably from one speaker to another. The ways in which vocabularies have evolved over time are recognized in the etymologies or word histories introducing each chapter.

In the stimulus control of nonverbal behavior, we say a response has generalized if a response maintained during one stimulus occurs when some new stimulus is presented. For example, if our tacting pigeon pecked the left key when a novel amber light was presented, we would say that the response controlled by red had generalized to the amber stimulus. A similar generalization of verbal responses to new stimuli occurs in the extended tact. Simile and metaphor provide familiar instances. We may say that someone is as busy as a bee or as sly as a fox (simile) or that someone is a hawk or a dove (metaphor). These extended tacts

presumably originated through generalization across some common properties sometimes shared by humans and bees and foxes and birds. It is through metaphor that language grows and changes (cf. Esper, 1918; Jaynes, 1976).

Another type of extension of the tact occurs when new words are formed by combining existing words (e.g., the words *type* and *writer* predated the invention of typing machines). Vocabularies change with changes in the environment important to the speakers of a language (but cf. Whorf, 1956). The ways in which tacts can be extended in human language are so varied that a detailed account is not feasible. Stewart (1975) has offered some interesting examples in his account of the origins of one particular class of tacts, place-names. For example, place-names are more likely to be based on unusual than on common features of a region. A valley in a forest of fir trees would probably not be named Fir Valley, but it might be named Oak Valley if a single oak tree stood there. Similarly, if wolves are common to a particular area, no stream there is likely to be called Wolf Creek, but a stream might be given that name if wolves are rare in the area and one has been sighted there.

Tacts of private events. Another important instance of the extension of the tact is to private events. Some tacted stimuli are accessible only to the speaker, as when we say we have a headache. Such tacts depend on the verbal community for their establishment and maintenance, but the problem for a behavior analysis is how the verbal community can establish and maintain these responses when it does not have access to the stimuli. A parent can teach a child color names because the parent can see the colors that the child sees and therefore can respond differentially to the child's correct and incorrect color naming. (So many different consequences follow from color naming that it ordinarily does not matter whether the parent teaches the color names through explicit instruction or simply allows the appropriate discriminations to become established through casual day-to-day interactions.) With private events, control can be established only through generalization from tacts based upon events to which the verbal community has access. For example, the child may learn to say when it is in pain because the parents have access to such overt manifestations as crying or facial expression, and if the child has learned the names of body parts, the two kinds of verbal responses may be extended to the tact of pain in a particular place.

But the verbal community's control of these tacts is weak, as demonstrated by the ambiguity of such responses when the overt manifestations are no longer available. For example, if someone says "I have a headache" and leaves a social gathering, it is not clear whether that verbal response tacted a private stimulus or simply allowed the speaker to escape from dull company. Even the language of private events depends on the public practices of the verbal community. The implications have been explored in considerable detail by Skinner and by Wittgenstein (see esp. Day, 1969, for a discussion of parallels between their treatments of private events). For our purposes, it is sufficient to note that verbal behavior does not ordinarily require that stimuli be simultaneously available to both speaker and listener. In fact, some of the important consequences of verbal behavior occur when the speaker can tact some event unavailable to the listener. Although the language of private events has its own particular difficulties, we need not invoke any new classes of verbal responding to deal with it.

Tact relations, by themselves, are only one part of verbal behavior, but through these relations verbal behavior comes into contact with the environment. Verbal behavior gains its power from these relations because without them there would be nothing of which we could speak. Some of what we call truth depends on the consistency of relations between verbal behavior and environment. Important correspondences between verbal behavior and the environment, however, depend not only on the tact relation, but also on the relations of verbal responses to each other.

Intraverbal behavior

Much of what we learn consists of verbal responses occasioned by other verbal responses. We learn to recite the alphabet, to count, to give appropriate answers to arithmetic problems, to recite poems, to provide definitions for terms, and to state facts. Utterances that seem superficially to have the character of tacts are often only verbal responses occasioned by other verbal responses. No one alive has seen Caesar at the Rubicon or Napoleon at Waterloo. We might argue that we have seen facsimiles, as in paintings, but even if we recognized this painted figure as Caesar and that as Napoleon, would we be likely to recognize the setting as the Rubicon or Waterloo unless the picture caption told us so? We do not ordinarily learn historical details in that way. Instead, given names or dates, we learn to say when or in what order events occurred.

In intraverbal behavior, one verbal stimulus sets the occasion for another verbal response. Free association provides a familiar example (Galton, 1879). The speaker, given a verbal stimulus, is asked to produce a verbal response. Because the immediate consequences of free associations are usually minimal, and because any given verbal stimulus may set the occasion for a variety of different responses, other variables may contribute to the particular verbal response given on any occasion. The procedure is therefore assumed to tap verbal responses that are of relatively high probability in the speaker's verbal repertoire.

In his discussion of intraverbal behavior, Skinner (1957) treats free association as follows:

One verbal response supplies the stimulus for another in a long series. The net effect is revealed in the classical word-association experiment. Here the subject is simply asked to respond verbally to a verbal stimulus, or to report aloud any responses he may "think of"—that is, find himself making silently. Echoic and textual responses are commonly produced but are either prevented by instruction or excluded from the results. Such an experiment, repeated on many subjects or on one subject many times, produces a fair sample of the responses under the control of a standard stimulus in a given verbal community. . . . The intraverbal relations in any adult repertoire are the result of hundreds of thousands of reinforcements under a great variety of inconsistent and often conflicting contingencies. Many different responses are brought under the control of a given stimulus word, and many different stimulus words are placed in control of a single response. For example, educational reinforcement sets up many different intraverbal operants involving the cardinal numbers. *Four* is part of the occasion for *five* in learning to

count, for *six* in learning to count by twos, for *one* in learning the value of π , and so on. On the other hand, many different verbal stimuli come to control the response *four*, e.g., *one, two, three. . .* or *two times two make. . .*. Many different connections between verbal responses and verbal stimuli are established when different passages are memorized and different "facts" acquired. The word-association experiment shows the results. (Skinner, 1957, pp. 73-74)

In its simplest form, intraverbal behavior has been the basis for a substantial part of the research on human verbal learning. The classic experiments of Ebbinghaus (1885) were specifically concerned with the learning of arbitrary verbal combinations. So that the learning would not be contaminated by previous language experience, Ebbinghaus created nonsense syllables for his experiments. In the current vocabulary, we would say that the experiments were designed to eliminate all other possible verbal relations but intraverbal ones. In paired-associates learning of nonsense syllables, one nonsense syllable is a stimulus item uniquely paired with another nonsense syllable as a response item. The learner's task is to give the appropriate response item for each stimulus item. We will not review human paired-associates learning here, because it will be treated in Chapter 12. It is sufficient to note that paired-associates learning represents a relatively pure case of intraverbal learning.

Another variety of intraverbal relation occurs in serial learning, the learning of a list of items in a particular order. Learning the arbitrary order of the letters of the alphabet and learning to count are common examples of serial learning. Such learning can create extended units of verbal behavior. Much research on intraverbal learning was based on the assumption that, in the production of word sequences, successive parts of the utterance served in turn as discriminative stimuli for later parts (cf. chapter 6 on chaining). Although the acquisition of sequential verbal responses may involve such relations (as in memorizing a poem or a definition or the order of elements in the periodic table), a point may be reached at which it is no longer appropriate to speak of each part of the utterance as providing the stimulus for the next part. Intraverbal control may continue as a component of extended utterances (it is most obvious if another speaker provides the stimuli, as when a prompter assists an actor by whispering the beginning of the next line in a script), but eventually an extended utterance itself may come to function as an independent verbal unit. The development of hierarchically organized, higher-order verbal units from verbal-response sequences is a critical feature of verbal behavior. Just as sequences of letters or phonemes are regarded as units when they occur in words, sequences of words may be regarded as units when they occur as phrases or sentences. Such maxims as "Haste makes waste" and "He who hesitates is lost" are reasonably viewed as unitary verbal responses. The analysis of such extended units and classes of such units in terms of their consistent sequential and relational properties is an analysis of behavioral structure.

It is paradoxical that, throughout much of the history of research on human verbal learning, the learning of intraverbal relations has been emphasized to the relative neglect of the learning of facts. One reason may be that it is simpler to define and manipulate verbal than nonverbal materials as stimuli. Another may be

that intraverbal relations are an important component of standard educational practice. For example, when a child who is mastering the multiplication table gives 42 as the response to 6×7 , the response is strictly intraverbal only if the correct response does not depend on other intervening arithmetic behavior (e.g., adding six sevens, counting by sixes, or counting the boxes in a six-by-seven rectangle). Before we can deal with the case in which the answer is derived rather than learned intraverbally, we must consider the consequences of verbal behavior and the effects of verbal instructions. We must also examine verbal interactions, and the behavior of the listener and reader. These topics prepare the way for a treatment of the complex manipulations in generating novel verbal behavior in novel situations.

The consequences of verbal behavior

We have emphasized the stimuli that control verbal behavior, but verbal behavior also has consequences. As with all operant behavior, these consequences may affect subsequent verbal responding. In a speech episode such as a simple two-person conversation, each person provides an audience for the other; each person, in other words, sets an occasion on which verbal behavior may have consequences. Audiences are varied in their stimulus properties; we speak into telephones, write messages, or address large groups of people. Often the consequences for the speaker are simply a listener's subsequent verbal behavior. It does not require a laboratory experiment to demonstrate that the listener's response can maintain the verbal behavior of a speaker. We tend to stop speaking to people who do not react to what we say. To this extent, we may say that the listener's responses reinforce the speaker's verbal behavior.

Some research on the operant control of verbal behavior has been on whether particular verbal responses, such as "yes" or "uh-huh," can reinforce a speaker's verbal behavior (e.g., Greenspoon, 1955; Azrin, Holz, Ulrich, and Goldiamond, 1961; Rosenfeld and Baer, 1970). Experiments have tested whether such consequences can modify the frequency of particular response classes (e.g., plural nouns) or the substantive content of a conversation. Results have been inconsistent, but failure to show that such consequences can affect the frequency of verbal classes does not bear on whether verbal behavior can be dealt with in terms of reinforcing consequences. Failure to demonstrate that a particular event serves as a reinforcer in a particular situation simply means that the term *reinforcer* is inappropriate to that instance.

The consequences that may serve as reinforcers for human verbal behavior, like the discriminative stimuli that set the occasion for it, are subtle and varied. The listener's responses may be verbal or nonverbal (e.g., answering a question, or coming to dinner when called), and responses that are effective in one situation or when made by one listener may be ineffective in other situations or when made by another listener. In addition, a particular consequence of verbal behavior may have different effects on the tendency to speak than on what is said. One difficulty with experiments on verbal reinforcers has sometimes been that they have attempted to modify the relative frequencies of different response classes with the same con-

sequences that are presumed to keep the speaker talking. In any case, it is important to consider what maintains both the speaker's behavior and the listener's behavior.

The mand as a class of verbal responses. One obvious kind of consequence for verbal behavior is illustrated when we are given something we ask for. If a child says "milk" and receives a glass of milk, we might say that the milk reinforces this verbal response. Verbal responses that specify their reinforcers have been called *mands* (demands and commands, for example, specify behavior in which the listener must engage). The response may or may not occur in the presence of the reinforcer. For example, a child may mand milk even if a glass of milk is not present. An analogue from animal research may be helpful. Assume that an experimental chamber for a rat contains one lever that produces food when pressed and another that produces water when pressed. If the rat presses the first lever only when food-deprived and the second only when water-deprived, we could argue that the presses are, respectively, food mands and water mands. Although it would not be essential, we could make the analogy more convincing by arranging signs that lit up for the experimenter when either lever was pressed, saying "Please give me food" and "Please give me water." Except that the rat's vocabulary is limited to two levers, the relations between the lever presses and their consequences are similar to those between verbal requests and their consequences.

Yet the account is not satisfactory. For example, we could imagine a child seeing a new toy, learning its name, and then asking for it, even though asking for this particular toy had never been reinforced in the past. We must conclude that, as a category of human verbal behavior, manding cannot consist simply of many separate response classes corresponding to each of the many potential consequences that could be manded (cf. Schick, 1971). Rather, manding is a single class of responses, characterized by the inclusion of the verbal response that in other circumstances tacts the reinforcing consequence. Within this general class, we can distinguish subclasses: e.g., mands that specify the delivery of particular stimuli ("Please give me an apple"), and mands that specify the behavior of the listener ("Please open the door for me"). In everyday discourse, we also distinguish mands by the consequences that may follow for the listener (e.g., pleas, requests, orders, etc.). For example, we speak of a demand when a speaker specifies an aversive consequence for a listener's noncompliance. Another important subclass consists of mands that specify the listener's verbal behavior (e.g., "Please give me your name," "What is this called?"). These mands may be further subdivided according to a variety of features. For example, we may speak of *prompts* when the appropriate verbal response is already known to the speaker (as in giving a hint when a child is unable to solve a riddle), and *probes* when it is not known (as in a police interrogation). We distinguish other properties of verbal mands when we speak of questions and instructions.

The role of instructions. We noted that verbal behavior can have either verbal or nonverbal consequences. In one way or another, the consequence is usually a change in the listener's behavior. For example, if we tell someone who is about to go outdoors that it is going to rain, we may change the likelihood that the person will take an umbrella. The distinction between production and comprehension

of language is mainly concerned with whether we emphasize the behavior of the speaker or of the listener. The verbal community maintains the correspondence between verbal behavior and environmental events in tact relations because only if this correspondence is consistent will the speaker's verbal behavior provide discriminative stimuli that can occasion appropriate behavior on the listener's part (Risley and Hart, 1968). If the speaker's verbal behavior is controlled by environmental events inaccessible to the listener, the speaker's verbal behavior may become a potent discriminative stimulus for the listener's behavior. For example, a listener's response to the tact, "fire," may have important consequences even if only the speaker has seen the conflagration.

Control by the speaker's verbal behavior will weaken if the speaker tacts unreliably. The probability of telling the truth or lying may depend on the potential immediate consequences of each for the speaker, but the long-term consequences of lying may be that listeners will stop acting on the basis of that speaker's verbal behavior (the tale of the boy who cried "wolf" illustrates these contingencies).

The listener's behavior may come to be strongly controlled by the major class of verbal behavior called instructions. A script is a set of instructions to an actor, and a text is a set of instructions to a reader. In both cases, the instruction specifies verbal behavior, i.e., what is to be said. The instructor who defines a term, for example, is specifying circumstances in which that term and its definition will be appropriate in the student's future verbal behavior. One important characteristic of instruction is that it often substitutes verbal discriminative stimuli for natural contingencies, as when a parent tells a child, "Don't touch the stove or you will burn yourself." This property of verbal instruction has far-reaching implications. In general, instructions may change the listener's behavior when the natural consequences by themselves may be ineffective or may be effective only slowly. If we invite friends to a party, for example, we give them directions rather than letting them search for the place on their own.

Consider a specific case. Someone without any explicit instruction types with the index fingers, one letter at a time, in the method called hunt-and-peck. For the novice typist, this method is faster than touch typing, in which each finger has a particular resting position at the keyboard. The immediate consequences of the two typing methods favor the former. But in learning touch typing, the long-term consequences of following instructions to place the fingers appropriately and to type particular letters with particular fingers eventually outweigh the short-term advantages of the hunt-and-peck system. What is learned is not solely a particular method of typing. To follow instructions successfully, the learner must ignore natural consequences (in this instance, the text is at first produced more slowly by touch typing than by hunt-and-peck typing). Because of the practical advantages of instruction, the verbal community shapes the behavior of following instructions across a substantial range of activities throughout a substantial portion of each individual's lifetime. Instructions become a class of verbal stimuli characterized by its effectiveness in overriding natural contingencies.

A major achievement of human verbal behavior is that it allows behavior to be controlled by descriptions of contingencies, in the verbal behavior of others,

rather than by the contingencies themselves. But the advantages of these unique properties of verbal behavior are accompanied by their special problems. A history of following instructions may make individuals susceptible to verbal control by authority figures (e.g., Milgram, 1963). In addition to the abuses that may arise when individuals simply follow orders, instructions may create problems in more subtle ways. For example, when a simple task such as pressing a telegraph key is established in humans by instructions instead of by shaping the key press, and key presses earn money according to various schedules of reinforcement, the performances of those given instructions are typically insensitive to the schedule contingencies compared to the performances of those whose responses were shaped (B. A. Matthews, Shimoff, Catania, and Sagvolden, 1977; cf. Reber, 1976). With uninstructed responding in humans, fixed-interval schedules produce their characteristic scalloping and ratio schedules produce higher response rates than yoked interval schedules (cf. chapter 8); neither of these schedule effects occurs reliably when responding is instructed. Just telling a human in an experimental setting to "press the key" will produce persistent responding that is insensitive to its consequences.

The problem is that we do not ordinarily want others to do what we say simply because we say it. A parent or teacher who gives instructions to a child to establish certain kinds of behavior might prefer but cannot be confident that the natural contingencies will eventually control the behavior and make instructions no longer necessary. For example, a reason for telling a child to put on overshoes before going out to play in the snow is that the instruction may avoid the later aversive consequences of the child's returning home with cold wet feet. But if the child always obeys the instruction, the natural contingencies will never have an opportunity to act on the child's behavior, whereas if the child disobeys the instructions, the aversive consequences that follow from snow and unprotected feet will enhance the control by instructions on future occasions. Thus, if we try to teach by telling others what to do, we may simply reduce the likelihood that they will learn from the consequences of their own behavior. There is no easy solution to this dilemma. We must always choose between the immediacy and convenience of verbal instructions and their longer-term effects on the learner's sensitivity to the consequences of behavior.

The listener's or reader's behavior. Given that the speaker's verbal behavior provides discriminative stimuli for the listener, the listener's behavior can be characterized simply as the behavior occasioned by these verbal stimuli. The listener's responses to verbal stimuli can be as varied as the responses to any other kinds of environmental events. Many of the possible verbal responses have already been considered, in the classification of echoic responses, intraverbal responses, and so on. Some nonverbal responses occasioned by verbal stimuli are also obvious enough that they do not require special consideration. Whether the critical stimulus is a red light, a traffic officer's outstretched hand, the utterance "stop," or a tree fallen across the road, the driver's response of stepping on the brakes illustrates the phenomenon of stimulus control. As we move from watching an actual incident to watching the incident acted in a play or a film and then to reading the script for the actual incident and then to reading a description of the incident in a story, the

common feature that holds these cases together must lie in consistencies of stimulus control over verbal and nonverbal behavior.

As with nonverbal stimuli, not all responses to verbal stimuli are operant. For example, if a spoken word is paired with a stimulus that elicits autonomic responses (e.g., electric shock), the word may come itself to elicit these responses. The phenomenon, sometimes called *semantic conditioning* (e.g., Riess, 1946), is a verbal equivalent to classical or respondent conditioning of nonverbal responses. Responding generated by these procedures generalizes across semantic as well as phonological dimensions of verbal stimuli. For example, if electric shock is paired with a farm word, such as *barn*, the conditioned galvanic skin response is more likely to generalize to other farm words, such as *cow*, than to words that simply have some letters in common with the original word, such as *burn*.

Yet if we say that a listener has understood something, it seems unlikely that we can provide an adequate account of the listener's response simply by appealing to the pairing of words with words (as in the giving of definitions) or of words with events (as in the teaching of tacts). The problem of meaning must reside at least in part in properties of the listener's responses to verbal stimuli. One critical property may be the correspondence between responses occasioned by a verbal stimulus and responses occasioned by the nonverbal events that the verbal stimulus ordinarily tacts. Many empirical studies of verbal behavior are primarily on ways in which verbal responses that tact related environmental events vary together in the verbal behavior of a speaker or have common effects on a listener. Whatever else is involved in the listener's behavior, it is reasonable to assume that the response to a tact shares some properties with the response to what is tacted.

An illustration comes from research on language in the chimpanzee (Premack, 1970), in which the chimpanzee's verbal responses, instead of speech, consisted of placing plastic chips of various shapes and colors on a magnetic board. Like sign language (cf. R. A. Gardner and Gardner, 1969), this language circumvents the limitations of the chimpanzee's vocal apparatus; in addition, its advantage over such transient verbal modes as speech or gestures is that it places fewer demands on memory capacity in establishing verbal behavior.

Some illustrations of the kinds of verbal stimuli and responses that a chimpanzee called Sarah could deal with are illustrated in Figure 10-3. Sarah was capable of varied verbal behavior, including not only manding and tacting but also the relational verbal responses we will later refer to as autoclitic. Here we consider only one example of the chimpanzee's verbal behavior. In a matching procedure, Sarah was taught to select from a pair of alternatives the property that matched a sample object. Given a round red apple, for example, she learned to select a red stimulus from a red-versus-green pair, and a round stimulus from a round-versus-rectangular pair. Sarah also made the same selections when the sample was the blue plastic chip that was the name for an apple. In other words, her response to the verbal stimulus, *apple*, was the same as her response to the nonverbal stimulus, an apple, that it tacted; the response was not controlled by the physical properties of the verbal stimulus itself (such transfer from a nonverbal stimulus to the corresponding tact may in fact be an appropriate criterion for calling responding verbal).




Figure 10-3. Some examples of verbal materials used in teaching language to the chimpanzee, Sarah. Each word consisted of a colored piece of plastic that could be placed on a magnetic board. Statements were formed by arranging the words vertically, as in the sequence on the left, *Sarah take apple*. The middle column shows abstract words: a negative, *no or not*; *name of*, eventually used to add words to Sarah's vocabulary; the interrogative, *?*, used in framing questions; and the comparative, *different* (other words in the abstract set included *yes*, *same*, and *if-then*). The top right, in which the scissors and the marble are actual objects, shows how the question, *scissors different from marble?* might be presented; Sarah's answer would be a choice of either the *yes* chip or the *no* chip. The bottom right shows the question, *name of apple?*, where the object on the right is an actual apple; Sarah's answer would involve replacing the interrogative chip with the triangular chip that is the name for apple. (From Premack, 1970, Figures 1-4.)

Human language also includes examples in which the properties that a verbal response ordinarily tacts override the properties of the verbal response itself. For example, if words are printed in different colors, it is difficult to tact these colors rapidly if the words themselves are incompatible color-names (e.g., the word *red* printed in green: Stroop, 1935); we read words, and do not ordinarily attend to physical properties such as the color in which they are printed. The relation between stimuli and the verbal responses they occasion clarifies some logical paradoxes of the language of reference. For example, if a self-descriptive word is homological (e.g., the homological word *short* is short), whereas a nonself-descriptive word is heterological (e.g., the heterological word *long* is not long), what then is the word *heterological*? That *heterological* cannot be homological without being a paradox of logic, but it is not a paradox of verbal behavior; a verbal response can tact a verbal stimulus, but it cannot tact itself.

In dealing with the formal relations, we saw that the symmetry and reversibility of stimulus and response relations increase the likelihood of speaking in terms of words rather than in terms of specific vocal or written modalities. Similar

correspondences exist in the relation between tacts and environmental events; these correspondences may be important when we speak of meaning independently of whether a word functions as a stimulus or a response. Consider, for example, rain as a stimulus, the word *rain*, and responses occasioned by rain, such as using an umbrella or putting on a raincoat. Both rain and the word *rain*, as stimuli, may occasion either a verbal response, the word *rain*, or a nonverbal response appropriate to rain. We may look out a window, see rain, and pick up an umbrella on the way out of the house, but the latter response might also be occasioned by a verbal stimulus, as when a weather report predicts rain. On seeing rain, we may speak of the rain to someone else, who may then take an umbrella, but this speaking might also be occasioned by the weather report. The relations are summarized in Figure 10-4.




Figure 10-4. Relations among verbal and nonverbal responses occasioned by verbal and nonverbal stimuli, where *S* = stimulus and *R* = response. Such consistencies are essential to understanding and meaning. (Cf. Figure 10-1.)

When we say that the listener, on repeating what the speaker had said, understood the speaker, we are ordinarily not satisfied to call the relation between *rain* as a verbal stimulus and *rain* as a verbal response echoic. Presumably then we judge understanding or meaning not by any single relation between stimuli and responses, but rather by the integrity of the sorts of interrelations illustrated in Figure 10-4. We may say that someone understands something that has been said when the individual can repeat what has been said not because someone else said it but for the same reasons that the other individual said it. Such behavior implies the consistent relations among verbal and nonverbal responses in Figure 10-4, and therefore must be central to the concepts of meaning and understanding.

In any case, these examples show that the response to a verbal stimulus may have properties in common with responses to the stimulus that is tacted. We know only a very little about the properties of these responses, but they are undoubtedly aspects of the phenomena we speak of in terms of meaning, reference, and semantics. Such behavior can be established by verbal instruction, as when we teach the names of things; the nature of instruction must therefore stand as a central issue in the analysis of verbal behavior.

*Autoclitic processes: Verbal
behavior conditional upon
verbal behavior*

We already noted that verbal behavior, like any other event, can be tacted. The functional relations that can be observed when nonverbal stimuli set the

occasion for verbal responses can also be observed with verbal stimuli. No new kinds of relations are involved, but the complexities generated when verbal behavior is built upon other verbal behavior need special comment. Verbal behavior that depends upon other verbal behavior for its occurrence and that modifies the effects of other verbal behavior is called *autoclitic*. It includes words that modify the effect of other words on the listener, and the complex orderings and arrangements that we call grammatical structure. A brief account cannot do justice to the richness of autoclitic processes; what follows will survey only a few representative examples.

Descriptive autoclitics. Many verbal responses tact the conditions under which other verbal behavior is emitted, and thereby modify the way in which the listener responds to what we say. In "I doubt the coffee is ready," "I think the coffee is ready," and "I am sure the coffee is ready," the speaker tacts some property of the tendency to say, "The coffee is ready," and the probable relation between that response and the condition of the coffee. The effect is to modify the likelihood that the listener will act upon the verbal response. The relation of *I doubt*, *I think*, and *I am sure* to the remainder of each sentence is analogous to that between the stimuli of a conditional discrimination. In "I recall that it rained yesterday," "I read that it rained yesterday," and "I heard that it rained yesterday," the speaker specifies the source of the verbal response, "It rained yesterday." Other descriptive autoclitics tact the speaker's reaction to current verbal behavior, as in "I am sorry to report. . ." or "I am pleased to say. . ."

Qualifying and quantifying autoclitics. We do not tact everything we see, and, conversely, we sometimes respond as if tacting when the stimulus is absent. The qualifying autoclitic accompanying such verbal behavior is typically some form of the verbal response, *no*. For example, we do not continuously say "The coffee is not ready" under conditions of unready coffee. Instead, this verbal response is likely to occur when circumstances set the occasion for saying "The coffee is ready" (e.g., the smell of coffee, or the question, "Is the coffee ready?"), but the conditions do not exist under which this verbal response can be uttered as a tact (i.e., the coffee is in fact not ready).

Assertion, like negation, is also autoclitic, but the verb *is* serves many functions. Sometimes it specifies that the verbal response it accompanies is a tact ("This is a book"), sometimes it prescribes equivalence between verbal responses ("A human is a featherless biped"), and sometimes, as already considered, it specifies temporal properties ("It is cold now"). The particular function of *is* often depends on other verbal responses or, in other words, on context. Not only does it function as a conditional stimulus with respect to the effect of other verbal behavior, but its function may in turn be conditional on other verbal behavior.

The effects of autoclitics may be quantitative as well as qualitative. Examples are *few*, *some*, and *many*, and the plural forms of nouns and verbs. The effect of *often* in "This text is often misunderstood" may be paraphrased as a statement that the tact, "This text is misunderstood," may often be uttered appropriately. As we shall see in the treatment of language structure, we have no independent nonverbal means for characterizing these relations, but paraphrase is often appropriate because it makes explicit the conditional relations among the components of the utterance. This is most obvious when the autoclitics are tacts specifying the

listener's verbal behavior (e.g., *vice versa* is conditional upon a preceding verbal response, and may be interpreted as a mand specifying that the listener generate a new verbal response reversing the order of particular components of the original verbal response).

Relational autoclitics. Some verbal responses can specify events only through their relations to other verbal responses. For example, words such as *above*, *before*, and *of* are not simply tacts of particular environmental circumstances. They occur in combination with other verbal responses, and depend on these other verbal responses for their effects. Because agreements of case and tense are also conditional upon properties of events tacted, these too can be treated as autoclitics. A structural unit of verbal behavior, such as a sentence in passive voice and past tense, can be regarded as a complex of tacts of various relational and temporal properties of a set of events.

An analogy from the behavior of a pigeon may be helpful. Suppose we add some new dimensions of stimulus control to the tacting of the earlier examples. We now project a circle or a square along with the colors in the window. We maintain left-key pecks during red, middle-key pecks during blue, and right-key pecks during green. But now we make reinforcement contingent upon fast pecking if a circle is in the window, and slow pecking if a square is there. Next, we vary the size of the circles and squares, and we make reinforcement contingent upon the force of the pecks: strong pecks when the figure is large, and weak pecks when the figure is small. If the pigeon's pecks come under the combined control of these stimulus dimensions, we will observe fast strong pecking on the middle key in the presence of a large circle on blue, and slow weak pecking on the right key in the presence of a small square on green. Just as a given set of events may occasion the separate words of a sentence in a particular grammatical order, the stimuli in the window may occasion responding that varies with their several properties. Incidental features of responding, such as the topography of successive pecks, may vary consistently with particular combinations of properties, just as stress patterns and other features of human verbal behavior may vary consistently across different grammatical forms.

We cannot say how difficult it would be to establish such a performance in the pigeon, nor whether the pigeon's response would be appropriate to a novel combination of stimulus properties (e.g., whether, if we had trained with respect to shape and size only during blue and green, the pigeon would peck with appropriate rate and force on the left key during red). The point is not to explore the pigeon's capacity for verbal behavior; rather, it is to illustrate the nature of the multiple dimensions of stimulus control over varied response properties in complex verbal behavior. To the extent that environmental properties occasion classes of verbal responses that are invariant even when they occur in combination with other response classes, we may treat such classes as verbal units (e.g., as when the present-tense active-voice sentence structure remains invariant across a variety of different tacted events). Such units may have complex structures, but it is only because such units and their components can be combined in novel ways that novel verbal behavior controlled by novel environmental conditions can be generated; things

previously unknown may come to be tacted only on the basis of verbal behavior that is available with respect to things known.

*Interaction between verbal
behavior and environment*

Our own verbal behavior may set the occasion for subsequent behavior. For example, we may act on a reminder we wrote for ourselves at an earlier time. Sometimes the subsequent behavior is verbal: we reword sentences, draw conclusions, derive solutions. These manipulations are of special interest when, as in problem solving, they lead to behavior that was unavailable earlier. The power of verbal behavior resides in how it provides discriminative stimuli setting the occasion for responding that may have important environmental consequences.

Consider an example from mathematics, a convenient illustration because the notation of mathematics is an exact prescription of the verbal responses appropriate to particular verbal stimuli. Understanding of addition or multiplication is judged by the number of ways in which an individual can respond to relevant verbal stimuli. An individual should be able to define these operations, to discriminate between cases to which they can apply and cases to which they cannot, to give answers to specific problems, and to derive each from simpler counting operations. Such behavior is verbal, and is necessary and sufficient for the statement that addition and multiplication are understood by that individual.

The learning of arithmetic involves intraverbal and autoclitic processes, but its particular advantages come when it is combined with tacting of the numerosity of environmental objects or events (cf. Ferster and Hammer, 1966). A child may calculate the number of objects in a rectangular array by multiplying the number of rows by the number of columns, or simply by counting all of the objects. Either operation is verbal, but the outcome is a verbal response that now may serve as a tact of the quantity of objects in the array. The structure of arithmetic corresponds to the structure of the environment so that new verbal responses occasioned by arithmetic behavior may then function effectively as tacts. This property of arithmetic behavior may be a property of verbal behavior in general. The components of thinking, deduction, inference, and other processes that have been called cognitive may not be as well defined, and may often be less accessible to an observer than the explicitly structured behavior of arithmetic. Nevertheless, they must share properties with the structure of the environment, because the environment is that of which we speak.

Psycholinguistics

SYNTAX: THE GRAMMATICAL STRUCTURE OF LANGUAGE

Sequential analyses

Constituents and phrase structure

Transformations

Hierarchical organization

Summary

SEMANTICS: THE MEANING OF VERBAL UNITS

The measurement of meaning

OVERVIEW

Language universals

Language development

ELEVEN

Three primary terms in the analysis of language have been *semantics*, the study of meaning, *syntax*, the study of grammatical structure, and *pragmatics*, the study of the functions of language. *Semantics* can be traced to the Greek *sema*, a sign or thing seen, *syntax* to the Greek *taxis*, arrangement, and *pragmatics* to the Greek *prassein*, to effect or do. *Pragmatics* is a relative of *practice*. *Grammar*, through the Greek *graphein*, to scratch or write, and *gramma*, a picture or a writing, is closely related to *graph*, *program*, and *topography*.

In this chapter, we turn from the functions of verbal behavior to its structure. Language is organized. Words are ordered in sentences, and sentences are ordered in texts. We also speak of how words are related to each other, as when we say that two words are similar or dissimilar in meaning. Psycholinguistics calls these topics *syntax* and *semantics*. The problem of *syntax* is how the words in a language are organized in sentences; syntax is therefore concerned with the *grammatical structure* of language. The grammatical ordering of words can be regarded as an autoclitic process. But in this chapter we will approach language from the structural point of view of psycholinguistics, noting the relation of this vocabulary to the functional account of chapter 10 when appropriate.

After syntax, we will consider *semantics*, which deals with the problem of meaning. When we introduced verbal behavior, we questioned the traditional concepts of meaning and reference. In this chapter, we will see what can be said about these traditional concepts. In effect, we will ask what it is that we tact when we say that particular words are related in meaning. Psycholinguistics deals with a speaker's vocabulary in terms of the speaker's *lexicon*, the dictionary of words available in the speaker's verbal behavior. We must therefore examine the structure of the lexicon. These topics, often related to each other, will pave the way

for our treatment of verbal learning and memory in later chapters. (Psycholinguistics has sometimes included a third topic, *pragmatics*, the uses of language; this corresponds most closely to the functions of language already considered in chapter 10.)

SYNTAX: THE GRAMMATICAL STRUCTURE OF LANGUAGE

In what ways can we describe how the words of a language are organized in sentences? We can count or transcribe or classify words with some consistency. We then can analyze sentences as orderings or arrangements of these discrete units. To begin, we know that we can discriminate grammatical from ungrammatical examples of the sentences of a language. Among grammatical sentences, we can classify or tact various categories of sentences, such as active voice, passive voice, past tense, and so on. But how do we define the dimensions along which we make these distinctions? We are concerned with describing the grammatical structure of language. This is a problem of stimulus control: sentences are stimuli, and their properties set the occasion for particular verbal responses, such as naming types of sentences. When we say that a sentence is declarative (a statement) or interrogative (a question), we are tacting something about the sentence. The grammar of sentences is primarily a structural issue. Before we are willing to say whether a sentence is grammatical, we do not ordinarily ask about the circumstances in which it was uttered or the effects it had.

We could begin by trying to list all grammatical sentences. But that task would be neither feasible nor relevant. Most of what we say or write is novel. Sentences are only occasionally repeated, and the possible lengths and varieties of sentences are unlimited. Thus, our list would be infinitely long. Even if we could produce one we would not have solved our problem, because we still would not know what made a given sentence grammatical. Instead of lists, we need an exhaustive description of the *kinds* of grammatical sentences. Although the number of possible sentences is unlimited, the number of kinds of sentences is presumably finite. In the language of psycholinguistics, we might speak of writing a grammar with a finite number of rules.

Three types of accounts describe the structural regularities of grammatical sentences (Chomsky, 1956; Chomsky and Miller, 1963; Catania, 1972). The first is restricted to *sequential dependencies* between successive words or word sequences, corresponding to the chaining of responses considered in chapter 6. The second describes grammatical structures in terms of the *constituents* or components of those structures. The third examines the *transformations* that show how a given grammatical structure with a particular arrangement of constituents is related to other grammatical structures. For example, consider the saying, "He who hesitates is lost." We would be examining sequential dependencies if we presented "He who hesitates" as a verbal stimulus and determined the frequency with which "is lost" is given as a response. If we instead studied how parts of the sentence are

named subject and predicate or pronouns and verbs, we would be dealing with constituents. And if we then described the relation of this saying to such paraphrases as "He is lost who hesitates," we would be dealing with grammatical transformation. We will find that each of these three structural features may contribute to our judgments of grammaticality. They are not mutually exclusive; none is adequate, by itself, to characterize the structural properties of verbal behavior.

Sequential analyses

We treated some properties of sequential analysis in the context of intraverbal behavior (chapter 10). There we considered the verbal responses to verbal stimuli called word associations. Word associations can be extended to verbal stimuli consisting of several words instead of a single word. Such sequences are the basis for constructing different *orders of approximation* to English. For example, suppose one is asked to complete a sentence given its first three words. The first of the three words is then dropped, and the two remaining words plus the next word are presented to a second participant, who now generates a sentence that continues those three words. Again the first word is dropped, and the next three remaining words are presented to still another participant, and so on. An example of a text produced in this way is:

the first list was posted on the bulletin he brought home a turkey will die on my rug is deep with snow and sleet are destructive and playful students always (G. A. Miller and Selfridge, 1950)

This is called a fourth-order approximation to English, because at each point a participant sees only the last three words of the sequence and must add the new fourth word. According to this nomenclature, a zero-order approximation to English has words chosen randomly, and a first-order approximation has words chosen randomly but in proportion to their frequencies in the language. The following compares samples from several orders of approximation to a sample of English text (from Miller and Selfridge, 1950):

Zero-order:	byway consequence handsomely financier bent flux cavalry swiftness weather-beaten extent;
First-order:	abilities with that beside I for waltz you the sewing;
Second-order:	was he went to the newspaper is in deep and;
Third-order:	tall and thin boy is a biped is the beat;
Fifth-order:	they saw the play Saturday and sat down beside him;
Text:	the history of California is largely that of a railroad.

Miller and Selfridge found that the more closely a text approximated English, the more accurately it was recalled. The critical point is that if we simply add new words on to the end of an existing sequence, in producing various approximations to English, we cannot guarantee that we will generate grammatical sequences.

Common sequences of words can occur together in ungrammatical ways ("Haste makes waste not want not") whereas rare sequences of words may be grammatical (e.g., "Sleeping green ideas dream furiously," "Eschewing obfuscation mitigates prolix propensities"). We can tact the grammaticality of novel sentences, but such tacts cannot be based on whether the word sequences were likely or unlikely. It will not help us much to examine sequential dependencies by recording the relative frequencies with which particular word sequences occur in some sample of sentences. In fact, we might judge some sequences as ungrammatical even though they are probable, as in the saying, "The wages of sin is death."

One reason that sequential analyses cannot adequately describe grammatical structure is that they fail to include relations among words independently of the separation of the words. Consider, for example, the relation between *chimpanzee* and *used* in the sentence "The chimpanzee used sign language." The same relation between these two words exists when the sentence is expanded: "The chimpanzee that was taught by the psychologist used sign language." Yet we cannot represent that relation with a sequential analysis. The probability with which *used* will follow *chimpanzee* in any sample of sentences is not likely to be the same as the probability that *used* will follow *psychologist*. Sequential probabilities when two words occur in immediate succession will necessarily differ from the probabilities when other words intervene.


Even if we could resolve that problem, we would encounter other difficulties. For example, sequential analyses would not help us to deal with ambiguous sentences. Consider the sentence, "Running experiments should be encouraged." This might be read as recommending more support for research ("Encourage the running of experiments") or more research on exercise ("Encourage experiments on running"). The grammatical structure depends on which interpretation is appropriate, but we cannot base our interpretation on a sequential analysis, because the same words occur in succession whichever way we read the sentence. The two structures therefore cannot be distinguished by word sequences. The circumstances in which the sentence was generated might be useful. For example, it might make a difference if the sentence were uttered in a psychology laboratory rather than at a track meet. But this is a functional distinction and not a structural one. In any case, the point remains that a sequential analysis is inadequate for dealing with sentence structure. This does not mean that sequential processes cannot or do not operate in verbal behavior. What we say at one moment must sometimes influence what we say in the next. We have called such relations intraverbal. We learn verbal sequences along with, or even independently of, grammatical structure (e.g., as when a child learns to recite the alphabet). Although sequential dependencies enter into verbal behavior, they cannot define grammatical structure.

Constituents and phrase structure


One aspect of grammatical structure derives from the constituents or components of particular sentences or classes of sentences. We determine the

constituents of a sentence by noting how its parts are related to each other. Consider the sentence "A word to the wise is sufficient." We can name its constituents: *word* is a noun, *is* is a verb, *sufficient* is an adjective, and so on. The relations are not defined simply by how close particular words are to each other. For example, the word *is* is tied grammatically more to *word* than to *wise*, even though it is closer to the latter. The relations among words in sentences have been represented in various ways (e.g., Wundt, 1900; Chomsky and Miller, 1963). Three types of representation, sometimes called phrase-structure diagrams, are illustrated in Figure 11-1.

I. BOX DIAGRAM



II. TREE DIAGRAM



III. BRACKETING DIAGRAM




Figure 11-1. Three methods for representing the constituent structure of sentences. Each shows how a sentence can be analyzed into structural units ranging from individual words to phrases. Note that the sentences in the box diagram (I) and the bracketing diagram (III) have equivalent structures. In recent psycholinguistics, the tree diagram (II) has probably been the most common form of representation. The places where the branches of the tree come together are called nodes; thus, the *verb phrase* in the example is a node at which *verb* and *noun phrase* come together.

We could go on now to discuss properties of noun phrases and verb phrases and so on. But the point is not to argue for one or another kind of description of the parts of a particular sentence. Rather, it is to illustrate the nature of an analysis

of constituents. The illustrations in Figure 11-1 mainly involve the naming or tacting of constituents. How do we decide whether a word should be called one type of constituent or another? We distinguish among constituents not by the words alone, but by their relations to other words in the sentence. We cannot classify constituents unless we also identify the structure of the sentence.

Consider the following example: "He whose laughs last laughs last." The word *laughs* appears twice, first as a plural noun and then as a verb; the word *last* also appears twice, first as a verb and then as an adverb. What we call *laughs* or *last* therefore depends on the relation of each word to other words in the sentence. This poses a problem. It appears that we cannot identify the structure of a sentence without classifying its constituents, but this last example demonstrates that neither can we classify the constituents without identifying its structure. How then can we ever describe the structure of a sentence? To resolve this problem, we must identify invariances of sentence structure. These invariances are not in particular words or word sequences, but rather in various coordinations among words. We can identify words of a sentence as members of particular classes (as particular kinds of constituents) because we have learned typical structures in our language (e.g., patterns of agreement between singular and plural nouns and verbs).

We may be tempted to look to the environment as a basis for the structure of sentences. But even if we show that certain kinds of sentences are likely to be uttered in certain kinds of circumstances, we cannot appeal to such circumstances as a basis for our naming constituents. This point can be demonstrated simply. We can name the constituents of some sentences even when the sentence is made up partly of nonsense words. For example, compare the sentences "He who guffs merts" and "She merts his guffs." Even though *guffs* and *merts* are not standard words in our language, we would call each a verb in the first sentence whereas we would call *merts* a verb and *guffs* a plural noun in the second sentence.

We must conclude that our grammatical classification of words does not depend on the environmental events of which we speak. The classification depends instead on the sentence structures within which words appear. Verbs, for example, are not defined as the class of words occasioned by activities; they are defined in terms of conjugation or other grammatical properties. Compare the activity word, *running*, in the sentences "The child is running" and "The child's running is fun to watch;" only the first instance is classified as a verb. The case is even more obvious when the activity word changes form with a change in grammatical structure, as in "The chess-player *moved* the knight" versus "The *move* surprised the other player," or "The jury *examined* the evidence" versus "Their *examination* was thorough."

It would be incorrect to assume that environmental determinants are never part of our judgments of grammatical structure. For example, in treating sequential dependencies we considered ambiguous sentences. An instance is "Time flies." In most circumstances, as when this sentence is uttered at a reunion of old friends, we call *time* a noun and *flies* a verb. But if the sentence were uttered in a biology laboratory as an instruction to record how long certain insects take to get from one place to another, we might conclude instead that *time* is a verb and *flies* is a plural noun.

We are left with a paradox. On the one hand there are sentences with structures that allow us to name the constituents without knowing the circumstances in which they were uttered (e.g., "He who gufs merts"); on the other hand there are sentences with structures that do not allow us to name the constituents unless we know the circumstances in which they were uttered (e.g., "Time flies"). Thus, any account of grammar that is either exclusively structural or exclusively functional will necessarily be incomplete.

We must understand what we have accomplished so far. We have examined some ways of talking about the structure of sentences, but we have not explained how we discriminate among nouns, verbs, and other parts of speech. We have not even discussed the dimensions upon which such discriminations might be based. We have merely shown that we can tact the constituents of sentences with some consistency (cf. Danks and Glucksberg, 1970; Gleitman and Gleitman, 1970, p. 10). So far, we have not identified any way independent of this tacting to specify the properties of constituents and their relations to each other. But this does not prevent us from speaking in the language of stimulus control. We faced the same problem in the treatment of abstraction. Properties of events set the occasions on which we respond with particular names, and verbal behavior itself can be included among such events.

In any case, and perhaps more important, our response to a sentence is typically not a matter of naming its constituents or drawing a diagram of its phrase structure. If someone asks us a question, we do not have to say which words are nouns and which verbs before we can answer. Children learn to speak and to understand sentences long before they are formally taught grammar and the parts of speech in school. Our analysis of grammatical structure therefore must be based on something more than our capacity to name constituents or to diagram structures. We must demonstrate the relation between these sentence properties and the speaker's or the listener's behavior. Experiments that seek to demonstrate such relations are said to be concerned with the "psychological reality" of various dimensions of language (e.g., Fodor and Bever, 1965; Schane, Tranel, and Lane, 1975).

Figure 11-2 provides an example (Johnson, 1965). People were asked to memorize sentences with different types of phrase structures. When they were later asked to recall the sentences, the probability of an error was greatest at the major grammatical transitions. For example, errors were highly likely at the break between the noun phrase and the verb phrase in both types of sentences, even though this break occurred at word transition 3 in sentences of type I and at word transition 5 in sentences of type II. The matching of the types of sentences in number of words and other properties demonstrated that grammatical structure and not simply numerical position or some other nongrammatical feature was the basis for the pattern of errors (cf. Jarvella, 1971, for a similar analysis at the level of phrase and sentence rather than word transitions).

In other experiments (Fodor and Bever, 1965; Bever, Lackner, and Kirk, 1969), listeners wearing earphones heard a sentence in one ear and a click in the other. They were then asked to locate where in the sentence the click occurred.




Figure 11-2. Probability of an error in sentence recall as a function of word transitions in two different sentence types. Sentences of type I are illustrated by: (The¹ tall² boy)³ (saved⁴ the⁵ dying⁶ woman). In this type of sentence, error probability was highest at the break between the noun phrase and the verb phrase (transition 3). Sentences of type II are illustrated by: [(The¹ house)² (across³ the⁴ street)]⁵ [burned⁶ down]. In this type of sentence, error probability was highest at the break within the noun phrase (transition 2), and the next highest at the break between the noun phrase and the verb phrase (transition 5). In other words, errors were more likely at those transitions corresponding to major breaks in phrase structure. (Adapted from Johnson, 1965, Table 2.)

When their judgments of click location were wrong, the clicks tended to be displaced in the direction of major sentence boundaries. This displacement did not depend on such acoustic properties as locations of pauses or speech intensity (Garrett, Bever, and Fodor, 1966). For example, the same tape recording was used to produce all but the opening words in pairs of sentences such as the following: (In her hope of marrying)^X(Anna^Y was surely impractical) and (Your hope of marrying)^XAnna^Y(was surely impractical). The parts of the recordings at which the major sentence boundaries occurred were acoustically identical. Nevertheless, a click in the middle of the word *Anna* was more likely to be heard displaced toward *X* than toward *Y* in the first sentence and toward *Y* than toward *X* in the second. These experiments have been controversial (e.g., Dalrymple-Alford, 1976; Glucksberg and Danks, 1975, chapter 4; Holmes and Forster, 1972; Reber, 1973; Reber and Anderson, 1970), on the basis of such factors as possible biases in the listeners' reports of sentence locations. Whether the results are interpreted as displacements in the heard locations of clicks or as systematic errors in the listeners' reports of those locations, they remain as demonstrations of the effects of sentence structure on the listener's behavior. This is the sense in which we may be justified in saying that such structures have psychological reality.

Transformations

In speaking about sentences, we may note that some sentences say different things while others say the same thing in different ways. In judging sentences this way, we are not discriminating among particular sentences or constituents. Instead, we are discriminating among certain relations between sentences. We cannot describe such relations solely in terms of constituents, because sentences with the same constituents in different orders may differ in grammatical structure. For example, consider the following three sentences: "He who hesitates laughs last," "He laughs last who hesitates," and "He who laughs last hesitates." The first two sentences have something in common that neither has in common with the last. We would say that the first two sentences "mean the same thing." Although the word orders differ in the first two sentences, the structural relations among their constituents are the same (e.g., *who* is similarly related to *hesitates* in both sentences). We therefore call the second sentence a grammatical transformation of the first.

Transformations describe changes in the structure of sentences that preserve certain relations among the constituents. When we transform a sentence from active to passive voice (e.g., "The rat pressed the lever" and "The lever was pressed by the rat"), we preserve the subject-object relations among the nouns and the verb. When we speak of different sentences in terms of the transformations that relate one to another, we are simply discriminating some of the structural features that the sentences share.

Once structural equivalences among sentences have been described by transformations, kinds of sentence structures may be distinguished by how they can be transformed. For example, consider the sentences, "He is hard to understand" and "He is last to understand." These sentences differ only in the adjectives, *hard* and *last*. They seem similar in structure, but we can demonstrate a difference by examining how each may be transformed. We can transform the first sentence to "To understand him is hard" but not to "He understands hard." Conversely, we can change the second sentence to "He understands last" but not to "To understand him is last." Although the original sentences appear similar, we can distinguish their structures by their different relations to other sentences. In the language of psycholinguistics, the sentences are called similar in *surface structure* (the particular order of constituents) but different in *deep structure* (the underlying structural features that distinguish them).

We sometimes judge how well a phrase or a sentence has been understood by whether it can be paraphrased appropriately (cf. Gleitman and Gleitman, 1970). For example, an answer is often related by grammatical transformation to the question that is asked. If a question like "Who bit the dog?" was answered by "The dog didn't bite anybody" instead of by "Nobody bit the dog," we might assume the question had been misunderstood. Again, we cannot appeal to semantics for our classification. The classification of transformations does not specify the conditions




Figure 11-3. Some transformations of the kernel sentence, "Haste makes waste." The middle block shows transformations from declarative to interrogative (statement to question, left-right), from active voice to passive voice (front-back), and from positive to negative (top-bottom). For each sentence type, the transformation to future tense is shown in the upper block and to past tense in the lower block. The closeness of the relation between two sentence types depends on the number of transformations that separate them (e.g., the future "Haste will make waste" is closer to the kernel "Haste makes waste" than to the past passive question "Was waste made by haste?")

under which any particular transformation occurs. Instead, it shows how different sentences may be related. For example, we can transform or paraphrase sentences even if nonsense syllables are substituted for some constituents (e.g., given the sentence, "The wizzle lerved the kalaka," we would identify "The kalaka was lerved by the wizzle," but not "The wizzle was lerved by the kalaka," as its equivalent in passive voice).

Figure 11-3 illustrates the relations among various transformations and combinations of transformations. The base sentence on which the transformations operate is "Haste makes waste." This base sentence provides the core material for transformations, and is called the *kernel* (in a stricter interpretation, even this present tense, active voice, positive, and declarative sentence is a transformation

from the raw material that makes up the sentence, i.e., the subject noun, "haste," the infinitive verb, "to make," and the object noun, "waste"). Possible transformations include present to past or future tense, active voice to passive voice, declarative to interrogative (question), positive (affirmative) to negative, or any combination of these. Even this list of transformations, which generates twenty-four unique sentences, is only partial (e.g., consider also changes from singular to plural, as in "The cat pauses" and "The cats pause," or changes in person, as in "I am bored," "He or she is bored," and "They are bored"). Furthermore, some transformations may be regarded either as simple transformations in their own right, or as complex transformations based upon a combination of simpler ones (e.g., the question "Does haste make waste?" might be a simple transformation of the kernel sentence "Haste makes waste," or it might be derived indirectly from the intermediate emphatic transformation, "Haste does make waste").

What advantage does speaking in terms of transformations have over simply labeling sentences according to tense, case, and so on? With transformations, as with constituent structure, we again face the problem of the behavioral significance (or psychological reality) of these properties of verbal behavior. One interpretation is that the transformations correspond to something a listener or reader actually does in responding to a spoken or written sentence. This is equivalent to saying that the transformation of a sentence from one grammatical form to another is a kind of behavior. Although such behavior may not ordinarily be observable directly, it might be measurable on the basis of such properties as its duration or its effects on other behavior (e.g., Mehler, 1963; G. A. Miller, 1962).

In one experiment (Savin and Perchonock, 1965), listeners were given sentences followed by sequences of eight words, as in the following: "The ball was hit by the boy. Grass, cat, truck, hour, desk, rain, shirt, blue." The transformational complexity of the sentences and the list of words varied from one presentation to the next. The listeners were instructed to repeat the sentence verbatim after each presentation and then to recall as many of the eight words as possible. On the average, the number of words recalled decreased with increasing transformational complexity of the sentences. For example, the mean number of words recalled was 5.27 after a kernel sentence, 4.67 after a question, 4.55 after a passive sentence, and 4.02 after a passive question. In other words, fewer words were recalled after sentences with two transformations (e.g., passive question) than after those with only one (e.g., passive), and fewer after sentences with only one transformation than after those with none (kernel sentences). The finding is consistent with the assumption that the transformational complexity of the sentence limits the listener's capacity to recall words from the list.

But other factors complicate the interpretation (Glucksberg and Danks, 1969; W. A. Matthews, 1968). Transformations are usually longer than kernel sentences (compare "The cat chased the mouse" with "The mouse was chased by the cat"). Even with sentences matched in number of words listeners typically take less time to repeat kernel sentences than transformations of the kernel. Thus, the results may depend on the time that repetitions of the various transformations


imposed between the end of the eight-word list and the opportunity to recall the eight words: the longer this time, the fewer the words correctly recalled.

Other experiments demonstrated that kernel sentences can be more quickly memorized than transformations (e.g., Mehler, 1963), and that grammatical structure affects the likelihood with which sentences are accurately heard when they are masked by noise (e.g., G. A. Miller and Isard, 1963). A more direct study of transformation as a kind of behavior gave readers two lists of sentences (G. A. Miller, 1962). Each sentence in one list was related by a particular transformation to one sentence in the other list. The reader was given one minute to match sentences from the first list to corresponding sentences in the second list. For example, sentence pairs such as "John warned the small boy" and "The small boy was warned by John" were used to make up two lists that differed by the transformation from kernel to passive. The time per transformation was estimated from the differences in the average time taken per correct match in various pairs of lists. According to this analysis, transformations from kernel to negative, from kernel to passive, and from passive to passive negative took from 1.1 to 1.8 seconds each, those from kernel to passive negative took 2.7 seconds, and those from negative to passive took 3.5 seconds. The different durations are consistent with the number of transformations by which the various pairs of lists differed: the first three pairs differed only by a single transformation, whereas the remaining pairs both differed by two transformations. (The transformation from kernel to passive negative can be broken down into two parts, first from kernel to either passive or negative and then to passive negative; the transformation from negative to passive is assumed to be indirect, by way of negative to kernel and then kernel to passive.)


These findings suggest that transformations, like constituents, have "psychological reality." We may recognize transformations as properties of verbal behavior, but we must not assume that transformations explain verbal behavior. We are able to describe some properties of combinations of transformations, as when we speak of a past-tense negative question in passive voice. We can note how transformations relate complex sentences to the simpler sentences from which they are derived (the psycholinguist speaks of the relation between a terminal sentence and its kernel or kernels). We might extend our account to more detailed features of grammatical structure (e.g., psycholinguists sometimes distinguish between obligatory and optional transformations, depending on whether general relations such as subject-predicate or other more variable properties of sentences are involved). Describing the structural properties of sentences is the point of grammar, but describing a sentence does not say how it was produced or understood nor does its production or understanding require a judgment about its grammaticality.

Hierarchical organization. Transformations can involve not only changes in the form of sentences, but also ways in which sentences or sentence parts can be combined with each other. For example, the sentences "The canary sang" and "The cat ate the canary" in combination can become "The canary that the cat ate sang" or "The cat ate the canary that sang." These types of sentence structures are called *recursive*, because the addition of new segments can recur again and again. With

the additional sentence, "The cat grinned," the combination sentences could become "The canary that the cat that grinned ate sang" or "The cat that ate the canary that sang grinned." Sentences can be expanded by adding phrases at the beginning or at the end, when their structures are respectively called left-recursive and right-recursive, or by adding parts within or around the sentence, when their structures are respectively called self-embedding and self-enveloping. With the self-embedded sentence "The canary that the cat ate sang," we can illustrate a self-enveloping structure by expanding it to "The music that the canary that the cat ate sang was off key." Whether we regard the structure of this sentence as self-embedding or self-enveloping depends in large part on whether we started with "The music was off key" or with "The cat ate." In the first case, the hierarchical structure might be represented as:



In the second case, it might be represented as:



In other words, we may describe a sentence as a hierarchical structure in which some features or components are necessarily subordinate to others. Whether a given component is higher or lower in such a description, however, may have nothing to do with the order in which components occurred when the sentence was produced or understood. For example, we might decide that one or the other of the above structures was an appropriate description depending on whether the sentence was part of a conversation that began with "I'm glad that the music stopped" or of one that began with "I'm glad that the cat ate." We may note again that we do not explain a sentence when we describe its structure.

Summary

We have examined some features of a formal analysis of syntax or grammatical structure. Sentences are sequences of discrete units called words, but sequential dependencies alone cannot be the basis for structural descriptions, because they cannot describe relations among words that are independent of the separation of the words in a sentence. Sentence structure, however, can be described in terms of constituents and transformations. In our account of a sentence, we classify words and their relations to each other, and we examine how the sentence is related to

other sentences that include the same or related words or word classes. The various methods of labelling or bracketing or diagramming sentence properties allow us to describe structural relations within and among sentences (as when we distinguish among the alternative structures of an ambiguous sentence). For any sequence of words, we must at least be able to classify its constituents and describe its transformational relations to other sequences. These properties of sentences may also have behavioral significance; to the extent that transformations describe what we do when speaking or listening or reading or writing, they correspond to the autoclitic processes discussed in chapter 10. Experiments concerned with the psychological reality of constituent and transformational structure illustrate how the behavioral effects of these dimensions of syntax can be studied.

SEMANTICS: THE MEANING OF VERBAL UNITS

To study semantics is to grapple with the problem of meaning. In everyday talk, we sometimes ask what people mean when they say certain things or whether one person has understood the meaning of someone else's words. A point early in our treatment of verbal behavior was that the traditional vocabularies of meaning and reference can be misleading. How can we reconcile the widespread appeal of the concept of meaning in our everyday language with the ambiguities of that concept in a behavioral account?

One part of the problem is distinguishing between the production of verbal behavior by a speaker or writer and its comprehension by a listener or reader (asking people if they mean what they say is different from asking them whether they understand what we mean). Even though correspondences between language production and language comprehension are ordinarily established in the course of language development, we must not assume that they necessarily go together. For example, a case report on an eight-year-old boy who could understand spoken English but was unable to speak because of a congenital organic defect (Lenneberg, 1962) demonstrates that comprehension can be independent of production. In the development of the psychology of language, behavioral accounts have tended to emphasize language production whereas cognitive and psycholinguistic accounts have tended to emphasize language comprehension. The treatment of semantics in this chapter, therefore, stresses comprehension: what happens when someone is said to understand a word or a sentence?

Let us begin with an experiment in which listeners heard various passages, such as an account of Galileo and the invention of the telescope (Sachs, 1967). After listening to the passage, they heard a sentence and were asked to say whether or not it had been in the passage. The sentence was either identical to one in the passage (base sentence) or differed from it in one of three ways: a change in word order that did not affect grammatical structure (formal change); a change in grammatical voice (active-passive change); or a change in meaning (semantic change). The sentence to be identified was presented either immediately after the

original sentence in the passage or after 80 or 160 syllables of additional text. Examples of the sentences used with the passage about Galileo are the following:

Base sentence—He sent a letter about it to Galileo, the great Italian scientist.

Formal change—He sent Galileo, the great Italian scientist, a letter about it.

Active-passive change—A letter about it was sent to Galileo, the great Italian scientist.

Semantic change—Galileo, the great Italian scientist, sent him a letter about it.

The passages and base sentences were varied across successive presentations and different listeners.

When the sentence to be judged immediately followed the text sentence, listeners identified this sentence as identical to or changed from the original sentence with better than 80% accuracy. When the sentence was presented after 80 or 160 syllables of intervening text, accuracy decreased for all sentence types, but accuracy for the semantically changed sentence remained greater than 75% whereas that for the other sentence types approached chance levels. In other words, the listeners were likely to recognize the sentence as different only if the meaning had been changed (semantic change); as long as the meaning remained the same, they did not discriminate the changed versions (formal and active-passive changes) from the original one (base sentence). The listeners were not remembering particular words or particular word orders. They were remembering something more fundamental: whatever it is that sentences have in common when we say they mean the same thing.

Other experiments also demonstrate that listeners are more likely to remember semantic structure than specific words or sentences (e.g., Anderson, 1974; Bartlett, 1932; Bransford, Barclay, and Franks, 1972; Fillenbaum, 1966; Kintsch, 1974). For example, in one study (Bransford and Franks, 1971) listeners heard a group of related sentences, such as "The ants were in the kitchen" and "The ants ate the sweet jelly." They later heard a selection from the original sentences plus new ones. Some of the new ones combined the original sentences, as in "The ants in the kitchen ate the sweet jelly." For each sentence, the listeners were asked to rate their confidence that they had heard it before. They were usually more confident that they had already heard the new combined sentences than they were that they had already heard the simpler sentences that had actually been presented. The listeners had apparently learned something more abstract than particular words or particular sentences (but see Katz and Gruenwald, 1974, and Wang, 1977, for alternative interpretations).

These findings should remind us of the intimate relation between semantics and syntax. When we examined transformations, we spoke of the psycholinguist's concept of deep structure; deep structure is the name for those properties of a sentence that remain constant over various transformations. For example, we say that "The teacher wrote the definition" and "The definition was written by the teacher" are the same or similar in meaning. This implies that what we hold constant when we change the syntactic structure of a sentence is its semantic structure.

The measurement of meaning

The experiments so far have mainly considered sentences, but more attention has usually been given to individual words. What determines what a particular word means? In what ways can the meanings of different words be related? Word associations were the basis for some attempts to measure meaning (Galton, 1879; cf. chapter 10). If listeners produced longer and more varied lists of words in response to one stimulus word than to another, the first word was said to be more meaningful than the second. In addition, it was assumed that the more closely related in meaning two words were, the more the lists of associates occasioned by the two words would overlap. For example, common associations to both "infant" and "baby" might include "crib," "bottle," and "diaper," but none of these might occur as responses to "guitar." The different degrees of overlap among the lists of associates are consistent with what we already know, i.e., that "infant" and "baby" are closer in meaning than either is to "guitar."

Some words occasion more associates than others. For example, concrete nouns like "wall" or "road" or "building" are likely to occasion more responses consistently than prepositions like "of" or "to" or "at." Given that responses may be occasioned even by nonsense words, the meaningfulness of such words relative to each other and to standard words in the vocabulary can be assessed (e.g., Glaze, 1928). Word associations are not equivalent to meanings, but to the extent that we say that words are similar in meaning when we respond to them similarly, such a treatment of word associations is appropriate as a measure of meaning.

There can be ambiguities in word association. For example, if the word "night" occasions the sequence "day, morning, sun, moon," should the later words in the sequence be regarded as responses to preceding words in the sequence or to the stimulus word itself? In this instance, we might guess that "sun" was probably occasioned more by "day" and "morning" than by "night," but whether "moon" was occasioned more by "sun" and "morning" (which shares three letters with it) than by "night" is unclear. The responses occasioned when each word in the sequence serves separately as a stimulus word might clarify such relations.

An attempt to measure meaning without such ambiguities was the *semantic differential* (Osgood, Suci, and Tannenbaum, 1957). Words were rated along dimensions provided by word pairs, like "happy-sad," "hard-soft," and "slow-fast." The similarities among the ratings of various words were then determined by a statistical procedure called factor analysis. This procedure created a space within which various words could be placed; the closer two words were to each other in this space, the more alike they were said to be in meaning. Words like "good," "beautiful," "clean," and "pleasant" tended to cluster together in this space, and to be distant from other clusters such as "bad," "foul," "dirty," and "ugly."

The semantic differential was intended to deal with all the words in the lexicon. More recent efforts based on word ratings have concentrated on words within specific categories, such as color names and the vocabulary of the emotions

(e.g., Fillenbaum and Rapoport, 1971). Still another technique, called *componential analysis* (Romney and D'Andrade, 1964), is based not on word ratings but rather on an analysis of the dimensions that distinguish words within particular word groups. For example, kinship terms in English are based upon sex (mother-father, sister-brother), generation (grandmother-mother, son-grandson), and line of descent (mother-aunt, son-nephew). The problem is that only certain word groups lend themselves to such an analysis (e.g., to what dimensions might the language of facial expressions be reduced?). Verbal classes with such ill-defined boundaries have been called fuzzy sets (e.g., Rosch, 1973).

Each method describes semantic relations among various words, and meaning might be defined simply as the relations measured by these methods, whatever their basis. This conclusion, however, seems unsatisfactory, perhaps partly because such definitions give us little to say about the behavioral significance of meanings. We already discussed meaning as a property of our responses to words. For example, if a familiar word like our own name is repeated over and over, we sometimes remark that the word loses its meaning; this implies that some response that occurs to the word at first drops out after several repetitions.

We have noted that the language of meaning can sometimes be misleading. For example, we are tempted to say in treating the preceding examples that we respond to the meaning of words. But meanings are not properties of the words themselves; they are properties of our responses to the words. A written word in a language that no one understands has no meaning. When a word is said to have many meanings (as in "bat" in a belfry and "bat" at a baseball game), it is the same word in its different contexts; the meaning of the word changes only in the sense that we respond to it differently. Consider saying that a red traffic light means "stop" and a green traffic light means "go." When we stop on red and go on green, we ordinarily say that we are responding to the red light or to the green light; we do not ordinarily say that we are responding to the meaning of the red light or to the meaning of the green light.

One implication of treating meaning as a feature of verbal behavior is that we may be able to measure some of its properties. For example, we might see whether some types of semantic behavior take longer than others (cf. the analogous earlier treatment of syntactic structures: G. A. Miller, 1962). In one experiment with this rationale (Collins and Quillian, 1969), readers were shown sentences and were asked to judge, by pressing one of two buttons, whether each sentence was true or false. The time taken to make each judgment was recorded. The true and false sentences were constructed from an assumed semantic structure consisting of a hierarchy of categories defined by sets of relevant properties. An example of such a hierarchy is shown in Figure 11-4. The assumption is that membership in the next higher class, both by a set of distinctive characteristics and membership in the properties "can sing" and "is yellow" and by membership in the class "bird." According to this account, the time taken for the judgment that a sentence is true should depend on where in the hierarchical structure a given characteristic is located. For example, of the sentences, "A canary can sing," "A canary has feathers," and "A canary breathes,"

the judgment should be quickest for the first sentence and slowest for the last. In the first case, the property "can sing" is a characteristic of canaries; in the second case, the property "has feathers" is a characteristic of being a bird, which in turn is a characteristic of canaries; in the third case the property "breathes" is a characteristic of being an animal, which in turn is a property of being a bird, which in turn is a property of canaries. (The structure allows for cases in which the characteristic property of a class may be incompatible with the properties that define the higher-level class of which the first class is a member; for example, because "can't fly" is a characteristic of penguins, judging the truth of "Penguins can't fly" does not depend on first judging whether penguins are birds.)

Judgment times were consistent with the types of hierarchical structure illustrated in Figure 11-4. Pressing the "true" button occurred a longer time after sentences like "A canary has skin" than after sentences like "A canary is yellow." But the hierarchical structures upon which the various sentences were based were constructed by the experimenters, and did not take into account such factors as the relative sizes of the different classes (e.g., Collins and Quillian, 1970; Kintsch, 1974; Landauer and Meyer, 1972; Oden, 1977; Wilkins, 1971). More important, later research (C. Conrad, 1972) showed that the different judgment times depended largely on the different frequencies with which particular word relations occurred in the language (e.g., "canary" and "yellow" and "sing" are more likely to occur together than "canary" and "breathe" and "skin"). The Collins and Quillian experiment stands mainly as a potentially useful methodology; it cannot be said that it has measured differences in semantic responses to words.

Yet our failure to provide an explicit definition of the word *meaning* does not imply that *meaning* is meaningless. The word is, after all, part of our verbal behavior. When we speak of what a word means, we are presumably tacting something about the relation between words and behavior. We may recall that an independent specification of what is tacted is not essential to the definition of a tact relation (cf. chapter 10). We have at least shown, however tentatively, that semantic structure can be described and measured (see also Kolers, 1966; Tulving,




Figure 11-4. Hypothetical semantic structure with three hierarchical levels. The classes at each level are characterized by sets of properties that define the members of each class. Only a few sample properties are shown for each class. (Adapted from Collins and Quillian, 1969, Figure 1.)

1962). We speak of certain classes of responses in terms of operants (chapter 6), and we speak of certain classes of stimuli in terms of concepts (chapter 7). But words can function either as responses or as stimuli. When we wish to deal with classes of events that function both as responses and as stimuli, it may be appropriate to speak of these classes in terms of meaning.

OVERVIEW

Children who grow up in China speak Chinese, and children who grow up in France speak French. Clearly they each have learned the specific grammar and vocabulary of their native languages. The question remains whether some general properties of human language that do not have to be learned underlie the details of particular languages. The extent to which the structural properties of human language are determined by biological constraints has been controversial (Lenneberg, 1967).

Language universals

Properties assumed to be characteristic of all human languages have been called *language universals*. Various grammatical relations have been proposed as universals. For example, it has been said that utterances in all human languages have subject-predicate structure, distinguish between singular and plural, and have two classes of constituents called open (expandable classes with no limit on the number of possible members, such as nouns and verbs) and closed (fixed classes with a limited number of members, such as prepositions and conjunctions). These properties may depend on the circumstances in which humans speak rather than on structural constraints on the kinds of sentences that humans can generate. The subject-predicate relation is one of our criteria for calling an utterance a sentence, and does not justify omitting utterances without subject or predicate from our analysis (e.g., "Hello," "Ouch!," "Oh?," and "Aha!"). The importance of distinguishing between single objects and collections of objects in all human environments must determine to some extent the distinction between singular and plural. And because some classes of words are semantically related to an unlimited variety of events (e.g., nouns and verbs) whereas others are syntactically related to a limited body of structural relations within sentences (e.g., prepositions and conjunctions), the respective classes may necessarily be open and closed.

The universals of human language have yet to be specified. The most convincing case for such universals would be the demonstration that certain types of sentence structures could not be learned, or at least that some types could be learned much more easily than others. But that humans might be predisposed to speak more in some ways than in others or that human speech might be constrained to develop in particular ways should not create problems for a behavioral account of language. Just as the different anatomical features of birds and bees and bats determine the different ways in which they each fly, the particular characteristics of humans may determine the structure of human language and the way in which it

develops. If we show that human verbal behavior always has certain structural properties or develops in certain ways, we have still to answer functional questions about the circumstances in which verbal behavior occurs (similarly, an anatomical account of the different properties of flight in birds and bees and bats would not tell us where or when one of these organisms would take off or land).

The issues have become more complex with demonstrations that chimpanzees can engage in verbal behavior. Early attempts to establish language in chimpanzees were unconvincing because they concentrated on language in the speech mode; the chimpanzee's capacity to engage in differentiated vocal behavior, however, is limited. Successful demonstrations of verbal behavior in chimpanzees came only after investigators switched to languages based on other types of verbal responses, such as gestures (R. A. Gardner and Gardner, 1969) or the manipulation of objects (Premack, 1970) or button-pressing at a computer console (Rumbaugh, Gill, von Glasersfeld, Warner, and Pisani, 1975). Debates then arose over the comparability of these and natural human languages. Once some features of language had been demonstrated in chimpanzee language, the feature could no longer be regarded as uniquely human (cf. Linden, 1976).

Language development

The nature of language development in humans can have important implications. For example, to demonstrate that some language skills differed from others in the extent to which they had to be explicitly taught presumably would affect educational practice. Most children, however, become fluent in their native language long before they begin formal education. Although there are many descriptions of the development of children's language with age, we have much to learn about the factors critical to that development (cf. R. Brown, 1973; McNeill, 1970; Segal, 1975).

An important feature of language is its novelty; when we write a sentence, it is likely to differ from any other sentence we have already written. We can deal with such novelty only in terms of the features that the new sentence shares with our earlier sentences; in other words, we must speak in terms of syntactic and semantic classes (cf. Esper, 1918). Such an account is appropriate even at the very early stages of language development.

An illustration is provided by Berko (1968). Sentence completions of children ranging in age from four to seven years were compared with those of adults. The children were shown and read an illustrated sequence of sentences that introduced a nonsense word. The sentence to be completed was designed to prompt a different grammatical form of the nonsense word. For example, a sequence that occasioned a plural form was: "This is a wug. Now there is another one. There are two of them. There are two _____." The regular plural in spoken English is formed by -z (as in "dogs"), by -s (as in "cats"), and by -es (as in "horses"), depending on the sound with which the spoken singular ends (e.g., -z after voiced endings such as *p*, *t*, *f*, and -s after unvoiced endings such as *b*, *d*, *v*). The children fairly consistently produced standard plurals with -z and -s (e.g., most children

responded *wugz* to the sample sequence above). Even though almost all of them were able to give *glasses* as the plural of *glass*, only about one-third gave the *-es* ending for the plurals of novel words like *tass* and *gutch* and *nizz*. Thus, various aspects of the standard plural usage in English appear to develop successively over time. Children may at first learn specific words and specific word sequences, but then they begin to master larger units such as particular classes of transformations (cf. Braine, 1963; Guess, 1969; Palermo and Eberhart, 1968; Palermo and Howe, 1970).

Other evidence that the acquisition of language is a progression from particular words to simple word groupings and then to syntactic and semantic structures of increasing complexity is in the usage of irregular verbs, such as "go," "come" or "break" (Slobin, 1971). Among young children, the regular but non-standard past tense forms are common: "goed," "comed," "broke." Children do not usually start with these regular forms, however. Instead, the child first learns the standard but irregular forms as individual words: "went," "came," "broke." Only after the child has learned some standard regular past tense forms are the irregular forms displaced by the regular but nonstandard ones, even though the irregular forms had been part of the child's vocabulary for some time. Months or years later, the standard irregular forms reappear to become permanent components of the child's verbal behavior. The progression from standard irregular to nonstandard regular and back to standard irregular past tense verb forms is consistent with a progression from the mastery of individual words through a stereotyped and inflexible syntactic form to the varied structure of fluent language.

These examples have been mainly of syntactic structure, but the same points apply to the development of semantic structure. A particular event in a novel context may occasion a novel grammatical utterance (as when a child gives the standard plural of a nonsense word), but the utterance may also be semantically new. We speak of such cases as examples of metaphor or analogy (cf. Jaynes, 1977; Tweney, 1974; and see chapter 10 on extended tacts). Metaphor and analogy can be effective only to the extent that there is some correspondence between the verbal relations discussed here in terms of semantic structure and the relations among events in the world. If we regard grammar as a description of the analogies that relate different syntactic structures, then we may also regard metaphor as the grammar of semantics.

Verbal Learning and Transfer

VERBAL LEARNING PROCEDURES

Serial learning

Paired associates

Free recall

Verbal discrimination

Verbal recognition: A special case of verbal discrimination

Summary

TRANSFER

Positive and negative transfer

Proaction and retroaction

Extensions to nonverbal procedures

Summary

TWELVE

Verbal learning procedures include *serial* and *paired-associates* learning, *free recall*, and the special case of verbal discrimination called *verbal recognition*. *Verbal transfer* and the phenomena of *proaction* and *retroaction* involve ways in which different verbal learning tasks affect each other. Many of these terms are closely related to words considered elsewhere: *associates* to *consequence*; *recall* to *class*; *recognition* to *cognition*; and *transfer* to *differentiation*. *Serial*, through the Latin *serere*, to arrange or attach, is related to *series*, *sort* and, perhaps as an arrangement of words, *sermon*. *Proaction* and *retroaction* combine the prefixes *pro-*, forward, and *retro-*, back, with a derivative of the Indo-European root *ag-*, to drive. This root is the source not only of the behavior word *act*, but also of such words as *agent* and *examine*.

The area traditionally called verbal learning is concerned with what happens as we learn word sequences, word combinations, and word contexts. These relations among words correspond to the intraverbal relations of chapter 10. In fact, the study of verbal learning might be defined as the analysis of intraverbal learning. This chapter examines the four major classes of verbal-learning procedures: serial learning, paired-associates learning, free recall, and verbal discrimination. The literature of verbal learning is extensive (e.g., McGeoch, 1942; Postman and Keppel, 1969; Tulving and Madigan, 1970); we will illustrate each class of procedures with a sample of some characteristic findings and some experimental and theoretical issues it has generated.

After this survey of procedures, the balance of the chapter will be devoted to some problems of transfer or learning: how does learning one set of verbal materials affect the learning of other materials? We will consider instances of both positive transfer (when the learning of the first set enhances that of the second) and negative transfer (when the learning of the first set retards that of the second). By then examining how current learning can interfere with what has been learned in the past and what may be learned in the future, we will extend our treatment to the phenomena called retroactive and proactive inhibition.

VERBAL LEARNING PROCEDURES

In a typical verbal learning procedure, we present verbal stimuli to the learner and then record verbal responses. Much is taken for granted in such arrangements. If we were to describe them simply in terms of verbal discriminative stimuli and differentiated verbal responses, we would be omitting some significant features. For example, the learner's performance is only rarely established by differential consequences; instead it is usually established by instructions. The learner who is instructed to learn particular items will ordinarily learn those items more quickly or more accurately than one not so instructed (this difference has been summarized by saying that intentional learning is superior to incidental learning; cf. J. G. Jenkins, 1933). Furthermore, the consequences of responding may be hard to specify. The experimenter designates responses as correct and incorrect or right and wrong. When the learner responds, the experimenter may say "right" or "wrong," or the learner may simply be shown the item designated as correct. These procedures are sometimes said to provide the learner with *feedback* or with *knowledge of results*. It is tempting to assume that telling or showing a learner that a response was correct is a reinforcing consequence, but such feedback can also be regarded as instruction (i.e., telling the learner that a response was correct is equivalent to saying, "respond the same way next time"). It may therefore be misleading to speak of the reinforcing properties of being correct or the punishing properties of being incorrect.

A variety of other circumstances may influence the learner's behavior. For example, it may make a difference whether the learner is a volunteer or serves in the experiment for hourly wages; if the experiment ends as soon as a list is learned, the consequences of finishing early will be different for the learner satisfying a course requirement than for the learner whose pay depends on the total time spent learning the list. These and other features of the experimental setting that determine the learner's performance are sometimes called *demand characteristics* (Orne, 1962). In verbal learning, we will mainly treat verbal stimuli and verbal responses; we will not be much concerned with the consequences of verbal learning. In some cases, it may be important to note such consequences, but usually we may assume that the consequences are fairly uniform throughout an experiment even if we cannot specify them.

Verbal stimuli and responses have ranged from such simple items as numbers or the letters of the alphabet through nonsense syllables or words to such complex materials as sentences or extended texts. Particular materials will be discussed in the context of particular experiments. Here we will only distinguish between *nominal stimuli* and *functional stimuli*. The nominal stimulus is the verbal item as defined or presented by the experimenter; the functional stimulus is that feature of the verbal item that occasions the learner's response (cf. attention: chapter 7). For example, if the stimulus items in a learning task are the four 3-letter sequences PDQ, TLC, VIP and ETC, but the learner's responses are occasioned only by the first letter of each sequence, then the three-letter sequences are the nominal stimuli and the first letters of each are the functional stimuli.

Verbal learning experiments have used verbal stimuli in both written and vocal modes, but written stimuli predominated in the early days of verbal learning research. The advantages of written stimuli are that they are uniform, easily described, and can be presented at well defined rates and durations. A popular device for presenting verbal stimuli was the memory drum, illustrated in Figure 12-1. Words were typed or printed on a paper loop or tape, and the roller advanced successive items to the window where they were presented to the learner at a controlled rate. In contemporary research, slide projectors or computer displays have replaced the memory drum. In early verbal learning research, vocal verbal stimuli were less favored because their uniformity and rate of presentation were harder to control. For example, an experimenter who had to read a list to several different learners might change the inflection or loudness of particular words over successive readings. The advent of tape recordings reduced the magnitude of this problem, and contemporary research uses either written or vocal verbal stimuli, depending on the suitability of one or the other to a particular experiment. Similarly, the choice of written or vocal responses varies with the requirements of particular experiments. Learners can usually respond more quickly vocally than in writing, but the vocal response has to be recorded whereas the written response is its own record.

The major classes of verbal learning procedures are summarized in Table 12-1. As its name implies, *serial learning* is the learning of a series or sequence of verbal items, as when one learns to count or to recite the months of the year. A series can be learned in a variety of ways. For example, suppose a memory drum is used to teach the first ten presidents of the United States. The drum could present each name for five seconds, and the learner could be asked to recite the entire list in proper order only after all ten names had been shown. This procedure, called *serial recall*, could then be repeated until the learner had given one or more correct repetitions.




Figure 12-1. A memory drum. Words on the paper tape are presented sequentially in the window for a specified duration. For example, in the serial learning of a 10-item list, the ten items might be repeatedly presented for 3 sec each. Depending on the specifics on the experiment, the learner might be instructed to anticipate each item before it appears, or to recite the list in its entirety at the end of each complete presentation.

TABLE 12-1

Classes of verbal learning procedures.

<i>Name</i>	<i>Description</i>	<i>Examples</i>
<i>Serial learning</i>	Learning the items of a verbal sequence in a specified order.	Learning to recite the alphabet; learning a poem.
<i>Paired associates</i>	Each of several different verbal stimuli occasions a particular verbal response. Each stimulus with its associated response constitutes a pair; the order of pairs is unspecified.	Given a state name, naming the state capitol; given a list of words in another language, providing the English equivalent for each.
<i>Free recall</i>	Naming the items in a list without regard to order.	Naming the members of the Supreme Court; naming the states.
<i>Verbal discrimination</i>	Items in one class of verbal stimuli occasion a different response than items not in the class; in other words, a discrimination in which the stimulus dimension is verbal.	Identifying the nouns in a paragraph of text; given names on index cards, sorting the cards into one stack of female names and another stack of male names.
<i>Verbal recognition</i>	A special case of verbal discrimination. The property that defines the class on which the discrimination is based is whether the verbal item appeared in a specified earlier list or context.	Given a list of names, saying which were from one's high school graduating class; distinguishing between old and new technical terms while reading a text.

titions of the list. Alternatively, as each name appeared in the window, the learner could be asked to say which name would appear next (e.g., on seeing *Jefferson* to say "Madison," and then on seeing *Madison* to say "Monroe," and so on). This procedure, called *serial anticipation*, could also be repeated until some number of correct repetitions. Once a list had been learned by either procedure, it might be asked whether the learner could answer various questions about the positions of names in the list, such as "Who was the ninth president?" (Harrison), "Which president was Tyler?" (tenth), or "Which president came before Van Buren?" (Jackson). The roles of context and position are illustrated by a common well-learned serial list, the alphabet. Even though we all can recite the alphabet and can probably say quickly which letter precedes or follows any other letter, few of us can give without counting the eleventh letter of the alphabet or the numerical position of the letter *S*.

In *paired-associates* learning, one learns the correspondence between items on two lists, as when one learns the dates of historical events or the equivalents in another language of particular English words. Each of several verbal stimuli sets the occasion for a different verbal response. For example, the pairs to be learned might consist of book titles and their authors. The sequence of items might appear in a memory-drum window as follows:

Don Quixote
Don Quixote—Cervantes
Candide
Candide—Voltaire
Moby Dick
Moby Dick—Melville
etc.

Each book title would appear alone for several seconds. Whatever the learner's response, the book title would then appear accompanied by the author's name. The full set of pairs would be repeated, perhaps with the order of titles changed, until for at least one repetition the learner correctly named the author of each book before the name was presented. Questions might then be asked about the effects of changing the response items or of reversing stimuli and responses. Experiments on paired-associates learning, however, do not ordinarily use items and pairings likely to be familiar to the learner; instead, they often use nonsense syllables or arbitrary combinations of items.

In *free recall*, the learner is asked to name, in any order, items that had been presented earlier, as when one is asked to say what was on a misplaced shopping list or what questions had been on a test. In the typical verbal learning study, the list consists of verbal items, but the procedure is like nonverbal recall, as when a witness names those who were present at the scene of a crime; whether the stimuli are verbal or nonverbal, the responses in both cases are verbal.

Verbal discrimination is simply discrimination along some verbal dimension of stimuli, as when a child is taught to distinguish vowels from consonants, or nouns from verbs, or grammatical sentences from ungrammatical sentences. As arranged with a memory drum, a verbal-discrimination procedure might present successive pairs of items in the window. The learner then must choose one item of each pair by naming it, pointing to it, or perhaps pressing a left or a right button. After the presentation of each pair, the learner is informed whether the choice was correct. Two examples of verbal-discrimination sequences follow with asterisks indicating the correct member:

JEG	VOB	JXF	MCF
	*	*	
ZID	FEP	QMH	DXJ
	*		*
BEW	DAX	FPW	ZJC
*			*
SEF	PIB	DHJ	ZGX
	*	*	
BUV	HIF	JFM	GZB
*		*	

In the sequence on the left, the designation of correct items is arbitrary; on the right, the correct item of each pair is the one with the letter *J*. An alternative procedure is to present items one at a time, with the learner indicating (e.g., by "yes" or "no") whether or not each item is in the designated class. As with other verbal learning procedures, presentations may continue until the learner meets some criterion of discriminated responding.

One type of verbal discrimination is that between new and old items, as when a student in a language course distinguishes words already learned from new vocabulary words. This is called *verbal recognition*. The learner is given an initial list. Later, the learner is given another list with both new items and items from the initial list. The learner must identify (recognize) the items from the initial list. Except that it uses verbal materials, verbal recognition shares much with familiar instances of recognition (e.g., as in recognizing friends in a group of people, or in recognizing places one has visited). All of these verbal-learning procedures have been applied to issues in the analysis of such areas as language and memory. We now consider each in more detail.

Serial learning

Verbal-learning procedures were an experimental realization of the associationist principles developed by such philosophers as David Hume and James Mill. The associationists had advocated that human thought was based upon the association of ideas. Ideas were said to become connected or associated in various ways (e.g., similarity, common elements, contrast), but especially through contiguity in space or time. Later, with the beginnings of modern chemistry, analogies were drawn between the formation of associations and the chemical combinations of atoms into molecules. Ebbinghaus (1885) saw the possibility of measuring the formation of arbitrary associations. He invented the nonsense syllable as an item that had not yet acquired verbal functions in the language and therefore would not be contaminated by associations that had already been established (but it was later shown that nonsense syllables did vary in meaningfulness as measured by word-association procedures: Glaze, 1928).

The typical nonsense syllable was a three-letter consonant-vowel-consonant sequence (a *CVC trigram*). Ebbinghaus constructed many such trigrams, excluding those that were already words in the language, and then arbitrarily made up lists of various lengths to be learned at a later date. Over a period of years, he learned the various lists he had constructed, recording such data as the number of repetitions required until he was able to reproduce a list without error. For a list of up to seven syllables, he required only a single reading for a correct reproduction; beyond that point, the repetitions required for a correct repetition increased with list length, up to about 55 repetitions for a 36-item list. Later research (cf. McGeoch, 1942) showed that not only total learning time but also the learning time per item increased with list length. Another of Ebbinghaus's findings was that meaningful material could be learned more rapidly than nonsense syllables. In contrast to the 55 repetitions required for a 36-item list of nonsense syllables, Ebbinghaus was able

to learn 80-syllable stanzas of poetry (from Byron's *Don Juan* in English) in about eight repetitions.

Serial-learning procedures, together with other verbal-learning procedures, provided the baseline for studying various factors that influence human learning. For example, the finding that spaced practice is typically more effective than massed practice (e.g., Underwood, 1961) is often cited as a basis for arguing that it is better to distribute study evenly throughout a semester than to cram it all in just before examinations. This finding is so familiar that it is sometimes difficult to recognize that early in its study it was regarded as counter-intuitive. It was argued that the massing of successive learning trials should allow less opportunity for the learner to forget items from one presentation to the next than if the trials were spaced in time. But this was not the case, and the superiority of spaced over massed practice eventually contributed to accounts of verbal learning in terms of interference between items (the superiority of spaced over massed practice, however, is much more evident in motor-skill than in verbal learning; Adams, 1954).

Other variables studied include meaningfulness (lists of words rated high in meaningfulness are learned more quickly than lists of words rated low), pronounceability (lists of items that are easy to pronounce are learned more quickly than lists of items that are hard), redundancy (lists of items organized in some way are learned more quickly than random lists), similarity (lists of dissimilar items are learned more quickly than lists of similar items, where similarity is measured by number of common letters or other shared elements), and so on (e.g., McGeoch, 1942; G. A. Miller, 1958; Underwood and Goad, 1951; Underwood and Schulz, 1960). A catalog of the factors that affect verbal learning (e.g., learning as a function of diet, or learning as a function of hours without sleep) can be expanded indefinitely, and the effects of any variable might depend critically on details of procedure.

One prominent feature of serial learning was the serial-position effect, illustrated in Figure 12-2 (Robinson and Brown, 1926). In a given list, items at the beginning tend to be learned most rapidly, and then items at the end, and finally items in the middle. The serial-position effect led to questions about how position in the list, as opposed to the preceding item, might occasion a particular response item. Serial learning had been regarded as a procedure in which each item was the stimulus that occasioned the next item. But the learner's responses can also be occasioned by questions about position (e.g., "What was the third item?" or "What was the next-to-last item?"). This finding led to attempts to define the functional stimuli in serial learning, i.e., those features of the list and the experimental situation that occasioned correct responding (Bewley, 1972; Woodward and Murdock, 1968).

How position can function as a stimulus, however, is not clear. For example, the beginning and the end of a list are usually marked by the pause that separates successive presentations. Such a pause is inevitable in serial recall, when the learner is instructed to reproduce the entire list after each presentation. In serial anticipation, on the other hand, the learner is instructed to name the next item before it appears, and the list can be repeated without interruption. In one experiment (Glanzer and Dolinsky, 1965), two groups learned a repeating 10-item list by the




Figure 12-2. Serial-position curves (percent correct as a function of position of the item) after 1, 5, 9, 13 or 17 presentations of a serial list. The data are averages across eleven learners, each of whom learned eight different ten-item lists of three-digit numbers. At all stages of learning, percent correct was lower in the middle of the list than at its beginning or its end. (Adapted from Robinson and Brown, 1926, Figure IX.)

method of serial anticipation (e.g., *ABCDEF GHIJ ABCDEF GHIJ ABC DE . . .*, etc., where each letter represents a different item). The first group was given standard instructions, but the second group was told that because of a procedural error they were starting with item *A* even though item *F* was actually the beginning of the list. The sequence of items was identical for the two groups, and yet the serial-position effect for the first group was appropriate to a list that began with item *A* (errors increasing with successive items after *A* and then decreasing toward the end of the list at *J*) whereas that for the second group was appropriate to a list that began with item *F* (errors increasing with successive items after *F* and then decreasing toward the end of the list at *E*). Because the lists were the same themselves; the data can be dealt with only in terms of the relation between the lists and the instructions.

Let us now return to Ebbinghaus and associationism. Specific issues arose from regarding serial learning as the formation of associations among items. Could associations be formed only between successive items, or were remote associations possible (e.g., between every other item, or every third item)? Were associations formed both forward and backward through the list? Ebbinghaus examined these questions by relearning lists of nonsense syllables one day after original learning,

and comparing the time taken to relearn each list when the order of the items was the same as or different from that of the original list. This was called the method of *savings*: the list was ordinarily learned more quickly on the second day than the first, and the effects of changing the order of items were measured by the different amounts of time saved on relearning. Ebbinghaus studied savings with 16-item lists in original, scrambled, and reverse orders, and with lists of every other, every third, every fourth, or every eighth item (e.g., in a list based on every third item, the order *A B C D E F G H* becomes *A D G B E H C F*). As shown in Figure 12-3, the greatest savings in relearning occurred with items in the original order, and the least with items in the scrambled order. The reversed list produced more savings than any of the remaining ordered lists. From these data, Ebbinghaus concluded that both remote and backward associations had been formed during the original learning of each list.

We have not defined associations, but let us assume for the moment that they simply correspond to discriminative control of a verbal response by a verbal stimulus. We might then say that once the list had been learned, a given item might occasion not only the next item but also the preceding one or the one after the next. If that is the case, serial learning may not be the way to study associations. It may be appropriate instead to examine associations independently of a particular serial order. Serial learning might then be regarded as a special case of paired-associates learning. For example, consider the serial learning of the following 5-item nonsense-




Figure 12-3. Original learning time (Day 1) for 16-item lists of nonsense syllables, and relearning time (Day 2) after various reorderings of the lists: no change; reverse order; orders based on every second, third, fourth, or eighth item; and scrambled order. The data for each type of reordering are averages across at least ten different lists learned on different occasions by Ebbinghaus. The shaded area shows the savings as a function of the different types of reordering. (Adapted from Ebbinghaus, 1885, Chap. IX.)

syllable list: NAJ, BEF, LUJ, PES, CED. In the presence of NAJ the learner is instructed to say BEF; in the presence of BEF, LUJ; in the presence of LUJ, PES; and in the presence of PES, CED. These same relations could be established in a paired-associates list consisting of NAJ-BEF, BEF-LUJ, LUJ-PES, and PES-CED. Presumably the learner who mastered the original serial list would quickly learn this paired-associates list, or vice versa (Postman, 1968). We therefore turn now to paired-associates learning.

Paired associates

Despite the relation between paired-associates and serial learning just suggested, paired-associates procedures appear to have been developed independently of Ebbinghaus's work on serial learning. The earliest experiments used colors rather than verbal items as stimuli; the verbal response items were two- or three-digit numbers (Calkins, 1894). Verbal materials for both stimulus items and response items later became fairly standard (Calkins, 1896), and paired-associates procedures, like serial-learning procedures, examined the effects of a variety of verbal properties (e.g., meaningfulness of stimulus items and response items: Cieutat, Stockwell, and Noble, 1958).

The learner's task appears simpler with a paired-associates list than with a serial list: each stimulus item occasions a unique response item, and the order of stimulus items is irrelevant. But even this relatively simpler task can be broken down into more fundamental components. Consider paired-associates learning as an instance of stimulus control. In contrast to stimulus control over a single response (e.g., as when a pigeon's key pecks are reinforced in green but not red), paired-associates learning involves many stimuli and many responses. Thus, the several stimuli must be discriminated, the several responses must be differentiated, and each response must be occasioned by its particular stimulus. Discrimination among the stimulus items is usually taken for granted; discriminations among various letter and word sequences have been well established in the human learner. But the learner must learn what the appropriate responses are before they can be occasioned by their respective stimuli. These components of paired-associates learning have been called *response integration* (or *response availability*) and *association* (Underwood and Schulz, 1960).

Consider the following three paired-associates lists:

DOG - XUW	DOG - 18	DOG - 3
ANT - WIX	ANT - 21	ANT - 1
CAT - VUW	CAT - 73	CAT - 4
FLY - XUV	FLY - 85	FLY - 2
BEE - XOY	BEE - 46	BEE - 6
COW - XIW	COW - 92	COW - 5

Each list has only six pairs of items, and therefore equal numbers of associations are to be established. With the first list, the learner would probably respond often with items not included in the response list (e.g., XIV, WUX, VIW), and might even have

difficulty with pronunciation. Less response integration would presumably be required with the second list, because of the learner's history with numbers. Nevertheless, the learner's responses would probably at first include numbers other than those on the response list (e.g., 72, 86, 93). We would say that response integration was complete when the learner's responses included all the numbers in the response list and no others. This could occur even while some stimulus items were not yet occasioning appropriate response items. Finally, the last list would require minimal response integration. Once the learner knew that the response list included only the digits 1 through 6, subsequent learning would comprise only associations, i.e., the particular relations between stimulus and response items. If paired-associates learning took much longer with the first list than with the last list, we might assume that the difference was mainly a matter of response integration or availability rather than one of association.

Even after it is separated from response integration, learning associations remains complicated. Presented by itself, a single stimulus-response pair presumably would be learned very quickly. To what extent, then, is the study of association complicated merely because many pairs are presented to be learned together at one time? (The circumstances are similar to those in the history of maze learning, in which mazes evolved through simpler and simpler forms, culminating in the *T*-maze and the straight-alley: cf. chapter 4.) These considerations led to studies on learning single associations, as opposed to learning many associations in a single setting (e.g., Estes, 1964). The issues rekindled a long-standing dispute in the psychology of learning, the *continuity-discontinuity* controversy. The question was whether learning occurred gradually and continuously, or took place discontinuously, in an all-or-none fashion. Were associations established a little at a time or all at once?

In one experiment (Rock, 1957), a learner was shown eight CVC-trigram pairs, and then was given an opportunity to respond to the eight stimulus items alone. This procedure was repeated until the learner was able to respond correctly to all eight stimulus items. In one group of learners, only the order of the pairs in the list was changed from one presentation to the next. In a second group, a pair was retained in the list only if the learner had responded correctly to the stimulus item; otherwise that pair was dropped and a new pair was substituted. In other words, for the second group, the only pairs retained over successive presentations of the list were those learned after one presentation. The two groups learned their eight-pair lists at roughly equal rates; learning was as rapid when new pairs were substituted as when pairs were repeated over successive presentations. Thus, trials on which a correct response had not yet occurred did not seem to contribute to learning.

Unfortunately, not all pairs of CVC trigrams are equally difficult to learn. For example, a nonsense syllable pair like CEN-TER, which forms a common word, will be more easily learned than a similar pair, like NEC-RET, that does not form a word. The problem with Rock's experiment was that the procedure for the second group was also a method for selecting a list of those pairs that were easiest to learn; pairs that could be learned on a single presentation were retained, and more difficult pairs were dropped. When the experiment was repeated with CVC-trigram pairs

that had been equated for difficulty, the group with a constant list learned more rapidly than the group for which new pairs were substituted for unlearned old pairs (Underwood, Rehula, and Keppel, 1962; see also Taylor and Irion, 1964).

We need not doubt that, at least sometimes, associations can be learned on a single trial. The reader who reads the pair OWL - NO, and then covers the pair up will certainly be able to respond appropriately later when the stimulus item, OWL - ?, is presented. Instead of asking how quickly the items of a single pair can be associated, it may be more important to ask about the properties of the association. For example, is it symmetrical? In other words, once the stimulus item consistently occasions the response item, will the response item be equally effective in occasioning the stimulus item (NO - ?)? The question is not as simple as it may seem, because failures to demonstrate symmetry may result from unavailability of the stimulus items as responses rather than from the reversibility of the association. (For example, the naming of written letters is a paired-associates task with written stimuli and vocal responses; if a child who does not yet write learns to say "A" when presented with a written A, we would not be surprised if the child could not write an A in response to the spoken letter.)

Even with such factors taken into account, associations are not symmetrical (e.g., Newman, 1972; Wollen, 1968). This finding has practical implications. For example, in learning another language, to be able to give the equivalent English word as a response to a word in the other language does not guarantee that one will be able to respond in that language when an English word is the stimulus item. The student is advised to learn the symmetry explicitly by studying the vocabulary in both directions (e.g., English-German and German-English). The existence of such symmetries is a defining characteristic of bilingual skill (e.g., Kollers, 1966).

Other practical findings follow from studies of paired-associates learning. For example, what is the most effective order in which to present pairs of items? Should easy items be presented before difficult ones? Once a stimulus item occasions a correct response, how soon and how often should the item be repeated? Experiments on such questions have led to methods of sequencing items that make paired-associates learning more effective than random sequencing or sequencing determined by the learner (Atkinson, 1972). Some other issues relevant to the effects of some item pairs on the learning of others are considered in the section on transfer of learning.

An account of paired-associates learning in terms of verbal stimuli that occasion verbal responses superficially seems adequate. But, as with serial learning, research has turned away from such an account. In serial learning, this was because verbal sequences can be formed into higher-order response units that are not reducible to sequences of stimuli and responses (e.g., in typing or writing the word *the*, the *h* is not simply a response to *t* as a stimulus: cf. chapter 6). Similar problems exist in paired-associates learning. Consider the following paired-associates list:

MON - CAR
GOL - BUT
FRI - CAB

SAT - YET
TAZ - BUS
ZIF - HOW

We might assume that the learner who masters this list has learned mainly associations. In such experiments, human learners usually do what they are instructed to do. But they usually do a lot of other things besides. The learner given the list above will probably note that only three stimulus items are nonsense syllables; the others are abbreviations for days of the week. Three response items are concrete, and might lead the learner to compose a sentence or a story that combines them. These and other factors may affect the rate at which the learner masters the list. We have emphasized the rote learning of arbitrary pairs of stimulus and response items, but paired-associates learning may depend on structural properties of lists that extend over many pairs of items, and may occasion behavior more varied than the designated response items.

By definition, an association has been learned when a verbal stimulus occasions its paired verbal response. Typically, such responding is established by instructions rather than by differential reinforcement, but that is not the reason why paired-associates learning is seldom treated in the language of stimulus control. The reason is instead that any human verbal learning that is not completely arbitrary will inevitably be more than the learning of associations.

Free recall

One demonstration that a learner who has mastered a paired-associates list has learned more than associations is that the learner can usually name some of the items even if no stimulus item is presented. For example, if asked about the last paired-associates list in the previous section, the reader might recall that the stimulus items included MON, FRI, and SAT. If the order of items is unimportant, this type of performance is called free recall.

Recall experiments can be conducted with nonverbal stimuli. For example, we could present a collection of objects and then, after its removal, ask someone to name the objects in the collection. A common instance of recall is when someone asks us to name the people we met at a party or some other social gathering. As a technical term, however, free recall is usually reserved for procedures with lists of verbal items.

In free recall, it becomes obvious that a major concern of verbal-learning procedures is memory. The defining characteristic of experiments on memory is that some time intervenes between a discriminative stimulus and an opportunity for the response that it occasions. Free-recall procedures usually present a list of items once, followed by an opportunity for the learner to name the items without the list. Thus, free recall resembles the first trial of a serial-learning procedure, except that the learner is not instructed to name the items in their original order. (Research has also examined recall of paired-associates lists, with variations called modified and modified modified free recall; Barnes and Underwood, 1959).

Given its similarity to serial-learning procedures, it should not be surprising that serial-position effects occur in free recall as well. Learners are most likely to recall items at the beginning and at the end of a list (Murdoch, 1962). These effects

are often described using two principles: primacy and recency. The principle of *primacy* states that the first items of a list are more likely to be recalled than later items; the principle of *recency* states that the most recent items (i.e., those at the end of the list) are more likely to be recalled than earlier items. It follows that items in the middle of the list are least likely to be recalled. But these principles do not explain serial-position effects in free recall; they merely summarize what learners do.

Primacy and recency can be separately experimentally (Glanzer, 1971; Glanzer and Cunitz, 1966). Groups of learners recalled items from a list immediately after the list was presented, after 10 sec of a counting task, or after 30 sec of a counting task. A primacy effect (relatively higher probabilities of recall for early items) occurred in all three conditions; the recency effect (relatively higher probabilities for later items) was evident with no delay, but was reduced or eliminated in the delay conditions. Imposing the delay prevented the end of the list from being the most recent event at the time of recall.

The position of an item in a list is only one of several factors that may influence its likelihood of recall. Consider the following 50-item list (after reading the list, the reader may wish to cover the page and try writing down as many of the words as possible):

DOG	SEAT	ELEPHANT	BEAR	TRACKS
TICKET	AIRPORT	CHIMES	PIPE	TIGER
RESERVATION	CABOOSE	WIGWAM	BISON	BUS
CAMEL	CHEROKEE	WOLF	FUTURE ~	PUEBLO
PLANE	WATCH	*destiny*	CHEYENNE	CLOCK
ORANGE	SQUAW	TRAIN	TIME	HIDE
MOHAWK	BOUNTY	FOX	SKIN	LION
FATE	PEEL	SEMINOLE	CHANCE	PAPOOSE
AUTOMOBILE	APACHE	HAND	BANANA	LOCOMOTIVE
APPLE	LEOPARD	GRAPE	CORN	MOCCASIN

Given the length of this list, the typical unpracticed learner will probably be able to recall considerably less than half the items. The earliest items (DOG, TICKET) might or might not be recalled, but their likelihood of recall will be greater than that of later items. Most learners will probably recall two or more of the items at the end of the list (LOCOMOTIVE, MOCCASIN). In addition, most learners will probably recall the distinctive item in the middle (*destiny*). The greater likelihood of recall of unusual items in a list is sometimes called the von Restorff effect (von Restorff, 1933). It can be demonstrated with a variety of distinctive features, such as color (e.g., a word in red on a page of standard print) or size (e.g., a word in large type). Semantic novelty may also be effective (e.g., a single common verb embedded in a list of animal names). The problem is that it is sometimes difficult to judge which features of an item will be distinctive in the context of a particular list. Suppose that recall probability is much greater for one item on a list than for neighboring items. We might then be tempted to say that the item must have been distinctive. But this does not explain anything; instead, it simply demonstrates one condition in which we call an item distinctive.

A drawback of attributing properties like distinctiveness or salience to stimulus items is in the implication that the important dimensions of free recall are in the list rather than in the learner. But the learner is not passive; the learner behaves with respect to the list. For example, in free recall learners typically rehearse recent items, by repeating them vocally or subvocally. In one procedure (Rundus, 1971; Rundus and Atkinson, 1970), items from 20-item lists of words were presented for 5 seconds each, and learners were instructed to rehearse aloud. Their rehearsals were tape-recorded. The learner had more opportunities to rehearse early items than later items (e.g., the fourth item could be rehearsed during presentations of any of the remaining 16 items, but the sixteenth item could be rehearsed only during presentations of the last 4 items). As illustrated in Figure 12-4, higher recall probability of early items (primacy effect) was correlated with the more frequent rehearsal of those items. The recency effect, however, was not correlated with rehearsal; recall probability increased for items at the end of the list, even though the opportunity to rehearse those items was limited. What the learner recalls may depend on what the learner does during and after the presentation of each item (perhaps a distinctive item is an item very likely to be rehearsed).

What a learner might do during free recall is not limited to rehearsal. We noted that the difference between serial recall and free recall is simply in the instructions to the learner. In serial recall, the learner is instructed to name the items in their original order; in free recall, the learner is not told that order is important, or perhaps is specifically instructed to ignore order. But this does not mean that items in free recall are named in arbitrary orders. In fact, the order of recall of items often differs systematically from the original order in the list. Consider the




Figure 12-4. Relation between number of rehearsals and probability of recall. Lists of twenty nouns were presented at a rate of five seconds per item. Learners were instructed to rehearse aloud, and were allowed two minutes of free recall after presentation of the list. (Adapted from Rundus and Atkinson, 1970, Figure 1.)

50-item list presented earlier. Words in various semantic categories are distributed throughout the list (e.g., American Indian words: MOHAWK, APACHE, WIGWAM, MOCCASIN, etc.; animal names: DOG, CAMEL, ELEPHANT, TIGER, etc.; travel words: TICKET, AIRPORT, TRACKS, BUS, etc.). At recall, the items within each category are likely to be named in clusters. The particular clusters may depend on properties of the list, current circumstances, and the learner's verbal history. For example, the recall of BISON in a cluster of American Indian words might occasion some of the other animal names; RESERVATION is more likely to be included in a cluster of travel words than in a cluster of American Indian words, because it comes early in the list near other travel words (TICKET, PLANE). A learner familiar with American Indian history or who closely followed American Indian affairs would probably recall more items in this category than someone without such interests; and so on. Occasionally, an item not on the list that is semantically related to words in a cluster is named; this type of error becomes more likely with longer lists and with delays between list presentation and recall (e.g., any of the following might appear in a recall list, even though only some are in the original list: NAVAJO, CHANCE, ENGINE, TEPEE, PANTHER, HIDE, TIGER, CAR, GIRAFFE, BUFFALO).

Clustering in free recall is most easily demonstrated with lists constructed to include semantic groupings (such clustering provides another method for studying semantic structure: cf. chapter 11), but it occurs with varied lists and along varied dimensions of verbal items (e.g., words that are common free-associates, or words related in spelling or pronunciation: cf. W. A. Bousfield, 1953; Brooks, 1968; Deese, 1959). The quantification of clustering can present problems, because consistencies of structure cannot easily be demonstrated in the single recall of a list by one learner. Techniques must be developed to cope with repeated recalls or with recall summarized over a group of learners (cf. A. K. Bousfield and Bousfield, 1966; Tulving, 1962). It remains clear that the learner is not passive when given verbal lists. The structuring of the items in a list during free recall is one kind of autoclitic verbal behavior.

Verbal discrimination

The most common verbal discrimination procedure is simultaneous discrimination: two verbal items are presented and the learner chooses one of them (by naming it, by pointing, or by pressing one of two buttons). The learner is then told whether or not the choice was correct. Unlike a simple discrimination studied in an animal laboratory (e.g., a rat's lever pressing reinforced in light but not dark), the procedure usually includes many different stimulus items in both the correct and the incorrect class. Verbal discrimination can also be established successively. For example, the learner might be given verbal items one at a time, and be instructed to respond "yes" or "no" to each item; alternatively, the learner might be instructed to sort a deck of cards with verbal items printed on them into two or more piles. Still another type of verbal discrimination is *visual search*, as when someone

is instructed to find verbal items in a list or in a text (e.g., Healy, 1976; Neisser, 1967, chapter 3). Examples are finding a name in a telephone book, looking up a word in a dictionary, and checking a manuscript for typographical errors. Experiments in visual search differ from the other types of verbal discrimination in that they are usually concerned with speed and accuracy rather than with the acquisition of control by stimulus items.

In simultaneous verbal discrimination, the procedure we will emphasize, the pairs of items and their assignment by the experimenter to the correct or incorrect class is usually arbitrary (e.g., pairs of nonsense syllables, with the items to be correct chosen randomly). With such lists, the acquisition of verbal discrimination depends on such factors as whether some items appear in the list more often than others (Ekstrand, Wallace, and Underwood, 1956). For example, a list in which each item appears only once is learned more slowly than a list in which some correct items are repeated.

Once verbal discrimination is established, one item in each pair occasions a particular response and the other does not. But the discrimination is arbitrary, and may not be relevant to the verbal discriminations already established in our everyday language (e.g., discriminations among nouns, verbs, and other parts of speech, or along semantic dimensions, or along dimensions of alphabetical order or spelling). Some experiments have explored the acquisition of verbal discriminations in which some feature is always a property of the correct item. For example, the correct item might be defined as an item that includes a particular letter or letter combination, or an item that does not include a vowel, or an item that falls into some semantic class. Such procedures are verbal cases of experiments on concept formation (cf. Figure 10-2).

In still another type of arrangement, whether an item is correct is determined by its context. Consider the following two verbal discrimination lists, in which each correct item is indicated by an asterisk:

DESK-TABLE
*
COUCH-BUREAU
*
CHAIR-DRESSER
*
CHAIR-DESK
*
TABLE-CABINET
*

WOOD-GLASS
*
SPOON-STEEL
*
RUBBER-KNIFE
*
CEMENT-NAPKIN
*
FORK-GLASS
*

In the first list, the designation of the correct item is arbitrary, and three items appear in two pairs. DESK is correct in both the first and fourth pairs, and TABLE is correct in neither the first nor the fifth pair. CHAIR, however, is correct in the third but not in the fourth pair. The learning of this list, therefore, requires a conditional discrimination with respect to this item: CHAIR is correct when paired with DRESSER but not when paired with DESK. In the second list, GLASS is repeated;

it is correct in the first pair but not in the fifth. The discrimination is based upon semantic classes: the correct items are all parts of a table setting, and the incorrect items are common materials. Because GLASS can be a member of either class, the learner's response must be based on the class membership of the other response. In this conditional discrimination, the relevant context is not only the item paired with GLASS but also the semantic structure of all of the pairs in the list. Including a conditional discrimination within a verbal discrimination list usually slows the rate at which the list is learned (Kausler and Boka, 1968).

A list with many conditional discriminations may be difficult to learn. Conditionality can also occur at the level of parts of an item. Consider the verbal discrimination lists that can be made from each of the following three sets of correct and incorrect items (correct items are marked by asterisks: cf. Chang and Shepard, 1964; Wickelgren, 1969):

*FIT	*FIT	*FIT
*FIN	*FIN	*FAN
*FAT	*PAT	*PIN
*FAN	*PAN	*PAT
PIT	FAT	FIN
PIN	FAN	FAT
PAT	PIT	PIT
PAN	PIN	PAN

From each set, we could make up a list of 32 pairs (each of the 4 correct items paired with each of the 4 incorrect items, and each of these 16 pairs with the correct item in the left or the right position). A verbal discrimination based on the first set would be acquired quickly; the learner would soon choose the word beginning with *F* on every trial. In the second list, no single letter is consistently in correct items, but some letter combinations appear only in correct items (FI-, PA-) and others appear only in incorrect items (FA-, PI-). But now note the properties of the third set. Every letter and every letter combination appears as often among correct items as among incorrect items (e.g., F-, -I-, or -T; PA-, P-N, -AN). In other words, given any part of an item, whether the item is correct is conditional upon the rest of the item; only the item as a whole can consistently occasion correct responses. Such relations between components of stimulus items and responses to these items are called *contingent associations*; lists constructed on such a basis can be exceedingly difficult to learn.

Verbal learning procedures usually repeat a list until the learner meets some criterion of correct responding. This is appropriate when the concern is with the relative difficulty of different types of lists. But when the concern is teaching, lists can be designed so that, through a progression of changes, they will efficiently produce particular verbal discriminations. In other words, verbal discriminations are well suited to fading procedures (cf. chapter 7). The following progression of items is an illustration (the correct item of each pair is indicated by an asterisk: adapted from Goldiamond, 1966):

A - *B
 *B - D
 *B - V
 O - *B
 FS - *BF
 *JB - EL
 *EAB - VOX
 *ABE - JILL
 IRENE - *BEN
 *BILL - CONNIE

DAISY MAE - *LI'L ABNER
 *ROBERT - MAY
 ANNA - *BOB
 *BOY - GIRL
 *DANNY BOY - BETTY ANN
 *BRUCE - BETSY
 RONA - *WILLIAM
 DOROTHY - *GEORGE
 *KENNETH - ANTONIA
 *JOHN - ELIZABETH

Early in the list, the discrimination is based simply on the letter *B*. At the transition from letters to words, the semantic difference, male-female, is introduced, with the male item always paired with the letter *B*. Later, the *B* is included in both items, and in subsequent pairs the *B* is removed so that responding is controlled solely by gender. The gender discrimination in the final pair also includes a reversal of the letter discrimination established at the beginning of the list. As a classroom demonstration, the fading from the letter to the gender dimension usually proceeds without errors after the first two or three pairs have been presented. In a more practical setting, the fading from pictures to written items may be an effective method for teaching early reading skills (Gleitman and Rozin, 1973).

*Verbal recognition: A special
 case of verbal discrimination*

The discrimination between items in a list may be based upon which item appeared in some earlier list. When this is the case, we speak of the discrimination as verbal recognition. For example, only one item in each of the following pairs appeared in the 50-item list used earlier to illustrate free recall: CHIMES-WHEAT, WIGWAM-TERMINAL, FUR-CORN, HIDE-BELL, ROOF-SEAT. The reader who correctly identifies the items from the original list is said to recognize those items. The response might be pointing at the item, naming it, or underlining it; in a less formal setting, the response might be saying, "I recognize this item as one that appeared on the list." This response is a tact controlled not only by the stimulus item but also by the circumstances in which that item was presented in the past.

The phenomenon of recognition is not restricted to verbal items. Discriminations based on whether people or places or objects are familiar or unfamiliar are examples of nonverbal recognition. One study (Shepard, 1967) compared recognition of words, sentences, and pictures. The procedure with words used a list equally divided between frequent words (e.g., *child*, *office*) and rare words (e.g., *julep*, *wattled*). Immediately after inspecting a 540-word list, observers correctly identified 88% of the original words in test pairs made up of words from the list and new words. Correct identifications were more likely for rare words than for frequent words. The procedure with sentences used a list of 612 sentences (e.g., "A dead dog is no use for hunting ducks"). Immediately after inspecting the sentences, observers correctly identified 89% of the original sentences in the test pairs made up of original and new sentences. In the procedure with pictures, recognition was ex-

amined at different times after observers had inspected a series of 612 pictures. After two hours, observers shown pairs made up of a picture from the series and a new picture correctly identified 99.7% of the original pictures. Correct identification decreased to 92% after 3 days, to 87% after 7 days, and to 58% after 120 days (50% is chance accuracy in this task). The study, however, does not permit a conclusion about the superiority of recognition for nonverbal visual stimuli relative to verbal stimuli. With pairs of pictures, the observer was discriminating between pictures seen before (from the original series) and pictures probably never seen before; with words (especially frequent ones), the observer was discriminating between words both of which had been seen before in many contexts but only one of which had appeared on the original list.

In general, verbal recognition does not require discriminations between novel and familiar stimuli; instead, it is likely to be based on whether the verbal item occurred in some previous setting (i.e., a particular list). A common error in verbal recognition is recognition failure (i.e., an item from the original list is not identified as old). False recognition (i.e., an error in which an item not from the original list is identified as old) usually occurs less often than recognition failure, and is most likely with new items semantically related to items on the list (cf. Anisfield and Knapp, 1968; Underwood and Freund, 1968).

Recognition may also depend on how new and old items are presented. When they are presented in pairs, the observer can make a correct identification by discriminating either the old item as old or the new item as new. But an old item can be presented *together* with several new items (e.g., as in a multiple-choice test). In recognition studies, the new items are sometimes called distractor items. Recognition accuracy decreases as the proportion of new or distractor items increases (Kintsch, 1968; Slamecka, 1967). If the observer guesses, the probability of a correct identification decreases with added distractor items (e.g., the probability of a correct guess is 50% when items are presented in pairs but only 25% when each old item is presented with three new items); the adjustment for guessing, however, is too small to account for the decrease in recognition accuracy accompanying an increase in distractor items.

It has often been assumed that recognition is involved in the performances established by other verbal learning procedures. For example, when a learner does not recognize an item as a stimulus item from a paired-associates list, that stimulus item will usually not occasion the appropriate response item (e.g., Bernbach, 1967; Martin, 1967). It has therefore been suggested that recognition of items as stimulus items from the list is a necessary condition for paired-associates learning. The trouble is that the dependency may also work in the opposite direction. The learner may sometimes recognize an item as a stimulus item from the list only because that item has occasioned a response item.

Verbal recognition requires discriminative responding in the presence of the item whereas free recall requires producing the item in its absence. We might expect these two classes of behavior to have different properties. Nevertheless, many accounts of verbal learning have assumed that recognition and recall are closely related. One account (Anderson and Bower, 1972) suggests that recall occurs in two stages: first the learner generates items, and then the learner produces only those

recognized as being from the list. Colloquially, we might say that the learner must think of the item first, and then must decide whether it was on the list.

One deduction from such an account is that recognition accuracy ought always to equal or exceed recall accuracy. In many studies, recognition is superior to recall even after the effect of guessing is taken into account (e.g., in a recognition test with paired new and old items, guesses will be correct half the time, but in recall of words or CVC trigrams, the probability of correct guesses will be small). There do exist circumstances, however, in which recall accuracy is consistently superior to recognition accuracy (Tulving, 1974; Watkins and Tulving, 1975). The most convincing cases are with cued recall rather than free recall; in this context, paired-associates learning is regarded as a type of recall in which the stimulus items help the learner to recall the response item of the list. For example, in one arrangement the learner was first given a paired-associates list (e.g., *glue-CHAIR*). Next, the learner was given free-association stimulus items likely to occasion response items from the paired-associates list (e.g., *table*, for which typical associates might be *CHAIR, CLOTH, DESK, DINNER*). Among the words produced in this free-association task, the learner was then asked to identify those that had appeared in the earlier paired-associates list (recognition task). Finally, the stimulus items of the paired-associates list were presented (e.g., *glue- ?*), and the learner's responses were recorded (cued-recall task). In procedures of this type, learners typically recognized fewer words from the original list than they were able to recall in the cued-recall task.

This outcome not only makes untenable the model of recognition as a component of recall; it is also likely to surprise those who regard recognition as somehow simpler than recall. The important point is perhaps that the comparison may not be appropriate in the first place. Recognition and recall are different classes of responses controlled by different classes of stimuli; these classes are measured in different ways, and therefore their comparison may not be justified.

Summary

We have examined the major classes of verbal-learning procedures: serial learning, paired-associates learning, free recall, verbal discrimination, and, as a special case of verbal discrimination, verbal recognition. In these instances of verbal learning, performances are typically established and maintained through instructions rather than through differential contingencies imposed on correct and incorrect responses. Although each procedure seems superficially simple, we discovered that the human learner, while mastering any of these tasks, inevitably does much more besides (in paired-associates learning, for example, the learner's behavior may come to depend on relations among stimulus items or response items across different pairs). We also concluded that it is inappropriate to regard verbal sequences or associations simply as stimulus-response units; human verbal behavior has other structural properties more important than the ordering of events in time (e.g., clustering in free recall). Now we turn to transfer of learning: once a performance has been established with one procedure, we may ask how the mastery of other procedures will be affected.

TRANSFER

Can the discipline of mathematics make one a more logical thinker? Will the mastery of a classical language like Greek or Latin make it easier to learn economics or history or sociology? Does the study of music or art establish skills that will be useful in studying great works of literature? Research on transfer of learning or transfer of training actually began with simpler questions than these (Thorndike and Woodworth, 1901). Some of the justification for such research was a common assumption: that the traditional educational curriculum was appropriate not simply because each subject matter might be important to the student in the future, but also because the disciplines "exercised the mind" or, in other words, taught various general intellectual skills. Studies of how learning one task could affect learning another were particularly relevant to this view. Yet each demonstration of the specificity of transfer to particular learning tasks or of some other limitation on transfer contradicted the view rather than supported it. And the typical college curriculum gradually evolved from a rigid sequence of required courses to electives, distribution requirements, and other flexible patterns of course options.

Thorndike may have played a critical role in this history (Joncich, 1968). His research interests gradually shifted from instrumental learning in animals to verbal learning in humans in educational settings; this change was paralleled by his increasing involvement with curricular matters such as course requirements for potential teachers. His findings in various studies of transfer were not only contributions to the psychology of learning; they influenced faculty committees responsible for establishing and reviewing course requirements and curriculum. The transfer of learning from one subject matter to another is difficult to analyze. It is reasonable to assume that students learn various study skills that generalize from one course to another. But to examine how the mastery of a classical language might influence a student's progress in mathematics or science or philosophy would require experimental procedures likely to be neither practical nor acceptable in the typical educational institution (e.g., random assignment of students to various sequences of courses; matching groups according to previous academic achievement and other criteria; controlling for differences in instructors and teaching methods in different classes; etc.). Perhaps inevitably, research on transfer concentrated on simpler tasks, such as transfer from one paired-associates list to another (e.g., McGeoch, 1942).

Positive and negative transfer

A complication in studies of transfer is that they necessarily compare learners. To say whether learning Task *X* affected the later learning of Task *Y*, we must know how the learning of Task *Y* would have proceeded if Task *X* had not been learned first. Task *Y* may be learned in a few trials or with few errors after learning Task *X*, but with only that information transfer from Task *X* to Task *Y* is indeterminate; Task *Y* may simply have been an easy task. Learning Task *Y* after Task *X* must be

measured relative to learning Task *Y* alone. When learning Task *Y* after Task *X* is essentially equal to learning Task *Y* alone, transfer is said to be zero (*zero transfer*: Task *X* did not affect Task *Y*). When learning Task *Y* after Task *X* is more rapid or with fewer errors than learning Task *Y* alone, transfer is said to be positive (*positive transfer*: Task *X* made Task *Y* easier to learn). When learning Task *Y* after Task *X* is slower or with more errors than learning Task *Y* alone, transfer is said to be negative (*negative transfer*: Task *X* made Task *Y* harder to learn).

This history of research on transfer includes many types of studies, both verbal and nonverbal (e.g., transfer of motor skills from one hand to the other: cf. Woodworth, 1938, chapter 8). Demonstration of positive or negative transfer in specific cases permitted only a cataloguing of particular transfer effects. Paired-associates learning suggested an analysis from which general principles of transfer might emerge; studies of transfer could be based on changes in stimulus items, changes in response items, or changes in their relations (Bruce, 1933). Table 12-2 illustrates several procedures in paired-associates transfer. The effect of learning one list on learning a second list can be examined when only the stimulus items are changed, when only the response items are changed, when both are changed, or when both remain the same but their pairing is changed.

TABLE 12-2

Some transfer procedures with paired-associates lists. In each hyphenated pair of letters, the first represents the list of stimulus items and the second represents the list of response items. Sample CVC nonsense-syllable items are also shown for each procedure.

Procedure	Designation		Sample items	
	List 1	List 2	List 1	List 2
Stimulus change only	<i>A-B</i>	<i>C-B</i>	lan-qip	fis-qip
Response change only	<i>A-B</i>	<i>A-C</i>	req-kiv	req-zam
Stimulus and response change	<i>A-B</i>	<i>C-D</i>	xal-pom	cam-lup
Same lists, stimulus and response items re-paired	<i>A-B</i>	<i>A-B_r</i>	hab-lef	hab-mot
			guv-mot	guv-lef

Changing only the stimulus items may produce positive transfer, especially if the original list included unfamiliar response items (cf. response integration in paired-associates learning); changing only the stimulus items, on the other hand, is likely to produce negative transfer, presumably because during the learning of the second list the stimulus items continue to occasion response items from the first list (e.g., Barnes and Underwood, 1959). Ideally, changing both the stimulus items and the response items should produce no transfer if the relations among items is arbitrary; in practice, the procedure assesses the effect of learning successive different lists (learning set, sometimes also called generalized or nonspecific transfer: cf. chapter 7). The re-pairing of unchanged lists of items is one of several other varia-

tions on paired-associates transfer procedures. For example, the symmetry of associations can be studied by exchanging stimulus and response items (transfer from an $A-B$ list to a $B-A$ list), and the role of particular properties of items can be studied by changing lists in various systemic ways (e.g., given A and A' as CVC nonsense-syllable lists differing only in item vowels, transfer from an $A-B$ list to an $A'-B$ list, or given B and B' as word lists in which corresponding items in each list are synonyms, transfer from an $A-B$ list to an $A-B'$ list).

Whether transfer from one paired-associates list to another is positive or negative depends on the relations among the items in the two lists. The relations have been summarized on the basis of the similarity between corresponding first-list and second-list items (Osgood, 1949). Because similarity can vary separately for stimulus items and for response items, the effects are represented as a transfer surface, illustrated in Figure 12-5. Positive transfer is maximal when the stimulus items




Figure 12-5. The transfer surface (after Osgood, 1949). The plane, $abdc$, represents the relation between list-1 items and list-2 items. Dimension ab represents similarity of stimulus items, and dimension bd represents similarity of response items. Points above the plane represent positive transfer (as at a , for $A-B$ to $A-B$ transfer), points below the plane represent negative transfer (as at c , for $A-B$ to $A-C$ transfer), and points on the plane represent zero transfer (as along the line from b , for $A-B$ to $C-B$ transfer, to d , for $A-B$ to $C-D$ transfer).

and the response items are identical in the two lists ($A-B$ to $A-B$ transfer); negative transfer is maximal when unchanged stimulus items must occasion second-list response items that are incompatible with the first-list response items ($A-B$ to $A-C$ transfer; re-pairing, in $A-B$ to $A-B'$ procedures, may also produce negative transfer). As the second-list stimulus items differ more from the first list stimulus items, transfer approaches zero ($A-B$ to $C-B$ transfer, or $A-B$ to $C-D$ transfer).

The empirical significance of the transfer surface, however, depends largely on the definition of similarity. Consider the learner who has learned LOUD as the response to a particular stimulus item. New lists might be constructed that substitute for LOUD response items related in spelling (CLOUD), related as synonyms (NOISY), or related as antonyms (SOFT). How are we to judge the similarity between List-1 and List-2 response items in these cases? Are words with similar spellings but different meanings more or less alike than words with different spellings but similar meanings? Are antonyms antagonistic response items, or should both antonyms and synonyms be treated as semantically similar? Should we call these examples of $A-B$ to $A-C$ transfer, or are the relations sufficiently close to justify speaking of $A-B$ to $A-B'$ transfer?

In fact, transfer would probably be positive in each case. The learner, having learned LOUD in response to a particular stimulus item in the first list, would readily learn CLOUD in response to the same stimulus item in a second list. But transfer from LOUD to NOISY or to SOFT would also occur quickly. Transfer might even be greater with the antonym relation than with the spelling or synonym relations. This finding might lead us to say that antonyms are semantically more similar than synonyms, and we might go on to argue that similarity is better measured in semantic terms than in terms of spelling or other properties. But sooner or later we would encounter problems. For example, spelling is likely to be a more important dimension of verbal items in some transfer procedures than in others. Although the transfer surface summarizes some features of transfer data, it is probably best regarded as part of the definition of similarity. Similarity is not a property of stimuli; rather it is derived from organisms' behavior with respect to stimuli. It follows that doing the experiment remains the most appropriate way to judge the amount and direction of transfer from one task to another.

An understanding of transfer effects may be critical to the sequencing of tasks in instruction. For example, transfer may help to establish some formal relations among spoken and written verbal stimuli and responses (chapter 10). Consider the formal stimulus and response relations as a child learns the alphabet:

echoic	S_{spoken} "a"	_____	R_{spoken} "a"
dictation	S_{spoken} "a"	_____	R_{written} A
transcription	S_{written} A	_____	R_{written} A
textual	S_{written} A	_____	R_{spoken} "a"

Ordinarily, echoic behavior is established long before a child learns to read and write. Thus, the transition to textual behavior might be regarded as an instance of $A-B$ to $C-B$ transfer (stimulus change only), and that to dictation might be regarded

as an instance of *A-B* to *A-C* transfer (response change only). We might guess that transfer in such cases would not necessarily be predictable from the transfer surface. For example, if a child had never learned to write letters, we would not expect much transfer from echoic behavior to dictation. If transcription had been established, however, we might expect that transfer to textual behavior would be easy. Thus, the order in which the formal classes are established may be a critical aspect of the teaching of reading and writing; determining the most effective order is a problem of transfer.

Proaction and retroaction

We have considered cases in which learning List 1 affects learning List 2. This is called a *proactive* effect, because the direction of effect is from an earlier to a later task. But learning List 2 may also affect subsequent List-1 performance. For example, after transfer from an *A-B* list to an *A-C* list, the learner who earlier had been able to give *B* response items to *A* stimulus items no longer may be able to do so (e.g., Barnes and Underwood, 1959); the *A* stimulus items now occasion *C* response items. This is called a *retroactive* effect, because the direction of effect is from a later to an earlier task. Most research on these effects has concentrated on cases in which the learning of one list interferes with rather than enhances performance on the other lists (Müller and Pilzecker, 1900); such effects are called *proactive* and *retroactive inhibition*, or *proactive* and *retroactive interference*.

As with other instances of transfer, proaction and retroaction are studied by comparing performance on a succession of tasks with performance on one task alone. The design of proaction studies can be represented as follows:

Experimental Group	Learn Task 1	Learn Task 2	Test Task 2
Control Group		Learn Task 2	Test Task 2

Except for designating part of the Task-2 performance as a test phase, this procedure is equivalent to the transfer procedures discussed earlier. The Task-2 learning is typically equated in some way for the two groups (e.g., by number of trials or some criterion of correct responding), and some period of time usually separates learning and testing of Task 2 (otherwise, if learning Task 2 continued to mastery for both groups, both groups would perform accurately on an immediate test and there would be no difference to observe).

The design of retroaction studies can be represented as follows:

Experimental Group	Learn Task 1	Learn Task 2	Test Task 1
Control Group	Learn Task 1		Test Task 1

For the experimental group, this procedure differs from the proaction procedure only in the task tested. For a given sequence of two tasks, therefore, Task 1 may have a proactive effect on Task 2 at the same time that Task 2 has a retroactive effect on Task 1. The two effects, however, would have to be assessed indepen-

dently in the two separate procedures. The procedures allow the study of proactive and retroactive effects among a variety of tasks (e.g., serial learning as Task 1 and paired-associates learning as Task 2, or French vocabulary words as Task 1 and Spanish vocabulary words as Task 2). In a broader context, what is learned in one course may be an important determinant of what is learned in a later course (proaction, as when preparation in algebra affects the mastery of calculus), but, conversely, what was learned in the earlier course may also be influenced by a later course (retroaction, as when the treatment of molecular bonding in an advanced chemistry course affects the terminology learned in introductory chemistry).

Extensions to nonverbal procedures

Instances of learning do not occur in isolation. Other learning preceded, and other learning is likely to follow. To the extent that every instance of learning may be affected by what came before and what follows, proaction and retroaction, and transfer in general, are relevant to all learning. A child's learning of the alphabet may be affected by earlier nonverbal learning of discriminations among geometrical patterns (e.g., Gibson, 1965). Upper-case letters, unlike lower-case letters, do not include any up-down or left-right reversals (e.g., *p* versus *b* and *b* versus *d*), and thus transfer from the upper-case to the lower-case alphabet may depend on whether the child has already learned nonverbal up-down and left-right discriminations. The verbal learning in reading a laboratory manual may affect the ease with which the student later masters various nonverbal laboratory skills. These examples are exceedingly complex. We have seen that even in simpler cases, such as paired-associates learning, the study of transfer requires an experimental analysis of various components of performance, such as association and response integration.

The simplest cases of transfer are those with a shift only from one discrimination to another. For example, the reinforcement of a pigeon's pecks in the presence of green but not red might be followed later by reinforcement in the presence of blue but not yellow or of circles but not squares. We might expect more rapid transfer from one color discrimination to another than from a color discrimination to a form discrimination. Similarly, reinforcement of a rat's jumps toward a left but not a right stimulus might be followed later by reinforcement of jumps toward a light but not a dark stimulus, or vice versa. We might find that transfer from position to brightness proceeded differently than transfer from brightness to position (the failure of transfer in such procedures has sometimes been called fixation: N.R.F. Maier, Glaser, and Klee, 1940).

Discrimination reversal may also be regarded as a type of transfer procedure. For example, the reinforcement of a rat's lever presses in the presence but not the absence of tone might be followed by reinforcement in the absence but not the presence of tone. The reversal of performance usually is more rapid if the stimuli are switched after many sessions of the original discrimination than if they are switched after only a few sessions of the original discrimination (the overlearning reversal effect or ORE: e.g., Capaldi and Stevenson, 1957; D'Amato and Jagoda, 1961; Reid, 1953).

Related phenomena are called *transposition* and *reversal* (e.g., Honig, 1962). For example, suppose a rat learns to respond to the larger of two squares in a discrimination apparatus. If the training squares are 2 and 4 centimeters in width and a new pair with widths of 4 and 8 centimeters is presented, the rat is likely to respond to the new larger square rather than to the now smaller square that it had responded to in the past. This transposition based on the relation between the stimuli will occur only within certain absolute limits; if much larger sizes are presented (e.g., 16 versus 32 centimeter squares), the rat may show reversal, i.e., responding to the smaller square, closer in absolute size to the one it had responded to in training. A theoretical account of transposition and reversal has been offered in terms of excitatory and inhibitory generalization gradients (Spence, 1937: cf. chapter 7).

Some cases of transfer change discriminations within the original stimulus dimensions (e.g., from one color discrimination to another), and are called *intra-dimensional shifts*. Reversals are one type of intradimensional shift. Other cases involve a change from one stimulus dimension to another (e.g., from color discrimination to form discrimination) and are called *extradimensional shifts*. The effects of stimulus dimensions in discriminative transfer have been assessed by comparing intradimensional reversal shifts with extradimensional nonreversal shifts (Kendler and Kendler, 1962). These are illustrated by the following two cases; each task is represented by sample pairs of verbal-discrimination items in which the first item of each pair is designated as correct:

<i>Reversal shift</i>		<i>Nonreversal shift</i>	
Task 1	Task 2	Task 1	Task 2
XON-map	MAP-xon	XON-map	MAP-xon
nij-TOY	toy-NIJ	nij-TOY	NIJ-toy

In each list, the items are CVC trigrams, but the item pairs differ along two dimensions: nonsense syllable versus common word, and upper-case versus lower-case. In Task 1 in both cases, the correct item is the nonsense syllable, whether in upper-case or lower-case. In Task 2 after the reversal shift, the correct item is now the common word; whether the item is upper-case or lower-case remains irrelevant. In Task 2 after the nonreversal shift, however, the correct item is the one in upper-case; the nonsense-syllable versus word dimension is no longer relevant. Analogous shifts can be arranged with nonverbal items (e.g., stimuli differing in both form and size). With adult humans, transfer usually occurs more rapidly with reversal than with nonreversal shifts, but the opposite is often the case with rats, monkeys, and very young children. These variations in performance have been interpreted in terms of changes in the control of responses by individual stimuli and by stimulus dimensions (e.g., Tighe, Glick, and Cole, 1971; cf. chapter 7 on attention).

Summary

Transfer among tasks is assessed by comparing learning a single task with learning two or more tasks in succession. Learning one task may enhance or inter-

ferre with learning other tasks (positive and negative transfer), and it may affect the performance on tasks learned earlier or later (retroaction and proaction). Whether transfer from one task to another is positive or negative depends upon the detailed relations of the stimuli and responses in the two tasks. The transfer surface appears to summarize these relations in terms of the similarity of stimuli and responses in the two tasks, but conclusions drawn from the transfer surface are made ambiguous by the nature of the similarity dimension; similarity is not a physical property of particular stimuli, but is derived from the properties of behavior in the presence of these stimuli. We have considered mainly examples of transfer from human verbal learning, but in closing we noted that the role of transfer is a general issue in learning, and we briefly considered some nonverbal examples of transfer. It is possible to make informed estimates of transfer in both verbal and nonverbal cases, but if data are unavailable the best way to find out about the dimension and magnitude of transfer from one task to another is to do the experiment.

Memory

THE NATURE OF MEMORY

MNEMONICS

THE VARIETIES OF MEMORY

Iconic memory: The persisting effects of stimuli

Short-term memory: The role of rehearsal

Long-term memory: Interference and the nature of forgetting

Overview: Metamemory as differentiated and discriminated remembering

THE METAPHOR OF STORAGE AND RETRIEVAL

Storage: Encoding and levels of processing

Retention: The question of memory reorganization

Retrieval: Cue-dependency and accessibility

THE STRUCTURE OF MEMORY

Episodic and semantic memory

SUMMARY

THIRTEEN

Memory has its source in an Indo-European root *smer-* or *mer-*, to remember, through which it is related to *remember* and *mourn*. The root does not seem to be linked to the Indo-European *men-*, to think, which is the source of *mnemonic*, *amnesia*, *memento*, *reminiscence*, *automatic* and, perhaps most interesting, *mind* and *mental*. *Forget* has a source in the Indo-European root *ghend-*, to seize or take. Through the Middle English *gessen*, to try to get, it is related to *guess*, and through the Latin *prehendere*, to hold before, it is related to *apprehend* and *comprehend*.

The term *memory* is the popular name for the topics treated in this chapter, but the chapter might as easily have been entitled "Remembering." The rationale for pre-ferring one usage to the other has been discussed by Woodworth (1921, pp. 5-6):

Instead of "memory," we should say "remembering"; instead of "thought" we should say "thinking," instead of "sensation" we should say "seeing, hearing," etc. But, like other learned branches, psychology is prone to transform its verbs into nouns. Then what happens? We forget that our nouns are merely substitutes for verbs, and go hunting for the *things* denoted by the nouns; but there are no such things, there are only the activities that we started with, seeing, remembering, and so on. . . . It is a safe rule, then, on encountering any menacing psychological noun, to strip off its linguistic mask and see what manner of activity lies behind.

The study of memory is concerned with ways in which an organism's responding now may be occasioned by some event or events in the past. A stimulus may set the occasion for responding even when a delay is imposed between the stimulus and the opportunity for a response. We will examine such situations, emphasizing what the organism does while the stimulus is present, during the time between stimulus and response, and when the response is occasioned. What the

organism does immediately upon presentation of the stimulus can affect whether the stimulus is remembered (or, in other words, whether the appropriate response occurs later). Accounts of memory often speak of what the organism does when the stimulus is presented as *memory storage*, and of what the organism does when the response is later occasioned as *retrieval* from memory. Systematic relations between stimuli to be remembered and subsequent responding occasioned by these stimuli are often discussed as *encoding*; for example, the learner who recites words aloud as they are presented visually is said to be encoding the visually presented words in an acoustic (vocal) modality.

Although most theories of memory have been concerned with how memory works, we will find that a dominant empirical theme in experiments on memory is the study of what is remembered. Analyses of what is remembered, like analyses of the effects of reinforcement, can tell us about important structural properties of operant classes and discriminative stimuli. Because studies of memory usually deal with the human learning of verbal material, they are especially relevant to determining syntactic and semantic structure. We will encounter this concern with the structure of what is remembered in the distinction between *semantic memory* and *episodic memory*. Semantic memory involves remembering various properties of the learner's language, such as word meanings; episodic memory involves remembering particular incidents at various times in the learner's past.

One strategy in studying memory is to examine the conditions under which organisms fail to remember. Important events can occur during any of the three stages of a memory procedure: stimulus presentation, imposed delay, and opportunity for a response. The effects of events occurring immediately upon presentation of the stimulus have led to a distinction between *short-term memory* and *long-term memory*. Stimuli may be quickly forgotten if the learner has to engage in other incompatible behavior immediately after their presentation and therefore is prevented from responding to them. The effects of events occurring during the delay between the stimulus presentation and the opportunity for a response have been discussed in terms of *interference*. It is in such cases that the language of *forgetting*, as the inverse of remembering, is most likely to be introduced. The learner may fail to remember if, during the delay, responses are occasioned that are incompatible with the response appropriate to the to-be-remembered stimulus. Accounts of memory in terms of interference are based heavily on empirical studies of transfer, and grow out of the analysis of negative transfer in studies of proaction and retroaction (cf. chapter 12). Finally, remembering may be affected by the conditions under which the learner is given an opportunity to respond. In general, the closer the circumstances are to those in which the stimulus was originally presented, the more likely it is that the learner will remember it. The study of *cue-dependent* or *state-dependent* learning or memory deals with such factors. The finding that particular stimuli may be remembered in some settings but not others leads to the distinction between *availability* and *accessibility* in memory: a stimulus that can be remembered is said to be available in memory, but it is said to be accessible only in those settings in which it is remembered; in those other settings in which it is not remembered, it is said to be available but inaccessible.

Memory is a complex subject matter with an extensive history (e.g., Murray, 1976). We discussed some of that history in the treatment of verbal learning (chapter 12), and will consider it further only in passing. We begin instead with two contrasting examples of memory. First we examine the language of memory in the context of relatively simple memory experiments with pigeons. Then we illustrate the complexity of human memory in the context of the memory strategies called *mnemonic* systems. These treatments set the stage for surveying various research on human memory.

THE NATURE OF MEMORY

In our analysis of verbal behavior (chapter 10), tacts were defined as verbal responses occasioned by nonverbal stimuli. The presence of the stimulus at the time of the verbal response was a constraint implicit in the definition of the tact. We may then ask how we should deal with verbal responses that depend on stimuli no longer present. For example, a witness to an accident may later describe the events at the scene, or may name the people who were present. These are the circumstances in which we might speak of the witness's memory or remembering (this is also the behavior called recall in chapter 12). But how much difference is there between tacting the events while they are occurring and describing them at a later time? The distinction is in the kinds of behavior that may occur between the event and its later recall. Such behavior may influence what is remembered.

To illustrate these points, let us return again to the pigeon. Our rudimentary example of tacting in the pigeon used a chamber in which pecks on a left or middle or right key were respectively reinforced during red or blue or green light in a window above the keys. In that example, the pigeon's pecks occurred in the presence of the lights. Suppose now that each light is presented only briefly, and the keys on which the pigeon pecks become available only after each presentation ends. If the red light appears for only one second and then the pigeon immediately pecks the left key, it would still be appropriate to say that the left-key peck is under the control of the red stimulus. We might still wish to call the left-key peck a tact of the brief red stimulus, even though that stimulus is gone when the peck occurs. Suppose we next impose a five-second delay between the brief presentation of red and the pigeon's opportunity to peck, and the pigeon continues to peck the red key. The left-key peck still is under the control of the red stimulus. We might be concerned with how the pigeon is able to respond consistently to the different colors after various delays, but we do not yet have a reason for assuming that the behavior is different in kind from that when the pecks occur in the presence of the colors. Imposing a delay between a stimulus and an opportunity for the response that it occasions does not necessarily alter the control of the response by that stimulus.

In fact, control of the pigeon's pecks by the color is likely to diminish even with a delay of one or two seconds, and responding will probably occur at chance levels at a delay of five seconds (cf. Blough, 1959). Yet we may be able to teach the

pigeon to remember. To do this, we might shape different performances after red and blue and green, so that each performance would later occasion a peck on a different key (e.g., pecking the left chamber wall after presentations of red, and pecking the left key if pecking the left chamber wall when the keys become available). If each color occasioned pecking in a different location, the duration over which the pigeon could remember a color would depend only on how long it could maintain its pecking in that location during the imposed delay. Under such conditions, control by the colors could extend over delays of many seconds and perhaps even minutes. By interrupting the behavior during the delay, we could then show how the pigeon's remembering depends on that behavior (Jans, 1976). For example, if we operated the feeder briefly between the stimulus presentation and the later availability of the keys, the pigeon would move to the feeder and eat, but it might not then return to pecking at the same location as before the feeder operation. In this instance, the interpolated feeder operation might reduce control of the key pecking by the colors to chance levels even at delays of only one or two seconds.

It is tempting to say that we cannot really call the pigeon's performance remembering if the temporal gap between stimulus and response is bridged by continued and localized behavior. Somehow, it seems appropriate to say that the pigeon's key pecks are more controlled by its previous behavior than by the colors as stimuli. Yet we modify our remembering by keeping calendars and appointment books. Consider the person who looks up a number in a telephone book and then begins reciting the number repeatedly until able to dial it. If the dialing is successful, we would probably be willing to say that the number was remembered, even though the dialing might have been more under the immediate control of the vocal repetitions than under the remote control of the number printed in the telephone book. Such behavior is sometimes called *rehearsal*, and it justifies the view that what is remembered is not so much stimuli as it is our own behavior toward those stimuli ("... what is reproduced on all occasions after the first is not the original but one's own reproduction of it": Zangwill, 1972, p. 130). These examples illustrate that questions about what memory is can often be translated into questions about what it is that we remember.

MNEMONICS

Memory may involve the persistence of behavior over time, as in rehearsal, but that is clearly only a small part. If someone remembers an event that took place yesterday or last week or last year, we do not assume that the individual was somehow rehearsing the event continuously from then until now. Mnemonics are techniques for increasing the likelihood of remembering (Bower, 1970). The use of mnemonic systems makes the learner less dependent on rehearsal, which is the major component of memorizing material by rote. A simple example of a mnemonic technique is the conversion of a sequence of symbols into a sentence, as when a beginning student of music remembers the notes on the lines of the treble clef, *EGBDF*, as the sentence "Every good boy does fine."

In order to form a series of places in memory, he says, a building is to be remembered, as spacious and varied a one as possible, the forecourt, the living room, bedrooms, and parlours, not omitting statues and other ornaments with which the rooms are decorated. The images by which the speech is to be remembered . . . are then placed in imagination on the places which have been memorised in the building. This done, as soon as the memory of the facts requires to be revived, all these places are visited in turn and the various deposits demanded of their custodians. We have to think of the ancient orator as moving in imagination through his memory building *whilst* he is making his speech, drawing from the memorised places the images he has placed on them. The method insures that the points are remembered in the right order, since the order is fixed by the sequence of places in the building. (Yates, 1966, p. 3)

For example, a student might imagine a systematic walk from one distinctive campus landmark to another: the dormitories to the dining hall, and next to the library, and next to the computer center, and next to the gymnasium, and so on. To learn some ordered series, the student then imagines each item at each successive location. To recall the series, the student would take the imaginary walk again, remembering each item in its appropriate place. To learn a new sequence at a later time, the student repeats the imaginary walk, in the same order, this time visualizing the new items in their successive places. The new series will be learned with relatively little interference from the first series (Christen and Bjork, 1976), but the first series may no longer be as well remembered. Thus, the method is useful mainly for series that need to be remembered only temporarily (e.g., a shopping list).

These mnemonic techniques work best with series of concrete items that are easily visualized or imagined. Unusual or bizarre items or combinations may have an advantage over common items, but even more important is the spatial closeness or connection of the items and places. For example, in learning pairs of objects, learners remember the pairings more accurately if the objects are presented in some relation to each other (e.g., one on top of the other) than if they are simply presented side by side (Wollen, Weber, and Lowry, 1972).

Various mnemonic devices were developed during Greek and Roman times, and in the Middle Ages became methods for remembering religious doctrine (Yates, 1966). They gradually became the basis for ritualized forms of religious art, as when particular figures were used to represent the vices and the virtues. In the course of this evolution, the mnemonic origins of these art forms were gradually

(and ironically) forgotten. Mnemonic techniques were given relatively little attention throughout much of the history of psychology, and major interest in their significance for analyzing memory has developed only recently (e.g., Bower, 1970).

Some techniques are designed so that abstract items can be converted into concrete ones. The *peg-word* technique is a simple method for translating numbers into a sequence that can be visualized, as in "one is a bun, two is a shoe, three is a tree, etc." An ordered list can then be learned by imagining each item together with the object corresponding to its numerical position in the list. These systems have the advantage over the method of loci that the learner can recall the item in any position without necessarily starting from the beginning of the list (e.g., to recall the third item, the learner has only to remember what was imagined with *tree*).

A more elaborate system provides a code for translating numbers into letters. This system has been a part of popular mnemonic techniques for roughly a century (e.g., Loiset, 1899). One version is the following:

Number	Consonants	Rationale
1	<i>t, d</i>	<i>t</i> has one downstroke.
2	<i>n</i>	<i>n</i> has two downstrokes.
3	<i>m</i>	<i>m</i> has three downstrokes.
4	<i>r</i>	<i>r</i> is the fourth letter of <i>four</i> .
5	<i>l</i>	<i>l</i> is the Roman numeral 50.
6	soft <i>g, j</i>	script <i>g</i> is an upside-down 6.
7	<i>k</i> , hard <i>c</i>	<i>k</i> can be combined with 7: <i>7k</i> .
8	<i>f, v</i>	both 8 and script <i>f</i> have two loops.
9	<i>p, b</i>	backwards <i>p</i> or upside-down <i>b</i> is 9.
0	<i>z, s</i>	<i>z</i> is the first letter of <i>zero</i> .

With this code, any number can be converted into a word or a sequence of words in which the sequence of consonant sounds corresponds to successive digits. With a little practice, the learner can quickly translate numerical information, such as dates or telephone numbers, into a form that is easily remembered. The potential applications are limited only by the learner's ingenuity. For example, consider the definition of the term *contingency*. This definition includes four key words, and the order of these words is critical: the *effect* of a *response* on the *probability* of a *stimulus*. The order can be remembered by translating their numerical positions into words made up of only the respective consonants: tea (1), Noah (2), may (3), and ray (4). *Tea* contains caffeine, which will have a drug *effect*; *Noah* built the ark in *response* to God's command; *may* implies a *probability*, as when we say that something may happen; and a *ray* of light is a kind of *stimulus*. This example is admittedly contrived, but it illustrates the flexibility of mnemonic systems.

At this point, the reader may wish to try an exercise. First choose a sequence of a dozen familiar locations, as in an imaginary stroll through one's neighborhood. The successive places should be distinctive, and ones that will be encountered in definite order. For example, they might include an apartment building, a playground, a school, a pharmacy, a parking lot, and so on. Next, construct an arbitrary sequence of twelve two-digit numbers (e.g., by consulting a random number table)

as, for example: 66, 57, 28, 14, 11, 87, etc. From the code presented above, each of these numbers can now be converted into a word. Each word should include two and only two consonant sounds; unsounded letters can be ignored. For example, 66 can be coded as "judge," since the "d" is silent. Concrete words are preferable to abstract ones (e.g., for 57, "lock" is preferable to "like"). Finally, each word, as derived, should be the basis for imagining an object or event at the successive locations (e.g., a giant *judge* sitting on the roof of the apartment building, a *padlock* attached to the playground swings, etc.). At some later time, the successive locations can be used to recall the words, and the words can be translated back to their respective numbers. On the first attempt with this system, it may be helpful to have the letter-to-number code on hand, but with some practice most users will find that an arbitrary number sequence can be learned with relatively little effort. Remembering fifty or more numbers is not a difficult task for a skilled user of mnemonic techniques.

Such arbitrary feats of memory might be useful mainly for impressing one's friends, but the main point is that they demonstrate the flexibility and capacity of human memory. Educational systems have tended to emphasize learning through understanding, and have correspondingly de-emphasized or even discouraged memorization. It is unlikely, however, that the learner will be disadvantaged by learning in more than one way, and therefore mnemonic techniques can be an effective supplement to other methods of study (cf. Higbee, 1977).

Mnemonic techniques work, but they are a far cry from the simple example of memory in the pigeon discussed earlier. The techniques illustrate how varied the phenomena are that must be considered in the study of memory. Even if we cannot say why humans are more likely to remember concrete items that have been imagined than abstract items that cannot easily be visualized, it at least is known that what the learner remembers depends on what the learner does. These mnemonic techniques are classes of behavior that can be learned.

THE VARIETIES OF MEMORY

The example of the pigeon's remembering required continued behavior that bridged the gap between stimulus presentation and a response after an imposed delay. The mnemonic examples did not include such uninterrupted intervening behavior. Clearly the cases are different. We would similarly distinguish between remembering an appointment because it is written down on one's calendar and remembering without consulting such a record. This implies that memory can take different forms. The period of time over which something can be remembered has often served as a basis for classifying the varieties of memory. We will examine three classes of memory phenomena: the relatively brief persistence of the effects of the stimulus; the maintenance of responding occasioned by the stimulus, as in rehearsal; and, appropriate responding after periods of time without rehearsal. These are often described respectively as *iconic memory*, *short-term memory*, and *long-term memory*, and have been the basis for many theories of remembering (e.g., Norman,

1968, 1970; Shiffrin and Atkinson, 1969; Sperling, 1963, 1967; Waugh and Norman, 1965).

Iconic memory: The persisting effects of stimuli

The effects of a stimulus may continue even after the stimulus presentation has ended. Persisting aftereffects of visual stimuli are called *icons*, and the topic of *iconic memory* is concerned with their time courses (for a discussion of the relation between icons and visual afterimages, see Sakitt, 1976). How is that time course to be measured? If an observer reports some letters that have been presented briefly, how is a report based on the sensory aftereffect to be distinguished from one based on the observer's continued rehearsal of the letters?

Humans are limited in the number of items that can be rehearsed or remembered after a single brief presentation. This limit, typically within the range of seven plus or minus two items (G. A. Miller, 1956), is called the *span of immediate memory*. For present purposes, it serves as a tool for studying iconic memory. The range of the span of immediate memory is to some extent independent of the nature of the items. For example, the learner will be able to remember roughly five to nine words almost as easily as five to nine letters, even though the words themselves include many more than nine letters.

The problem can be solved by showing more items in the stimulus display than can be included within the span of immediate memory and asking the observer for some sample of these items at various times after the display. Just as an examination tests only part of the material that a student is supposed to have learned during a course, and assumes that the score represents what proportion of the entire course has actually been learned, the sampling assumes that the observer's responses represent the proportion of items that can be reported at various times after the display. The time course of iconic memory is measured over fractions of a second, but during this time the observer in effect can still read the items in the display even though the display is no longer there.

This experimental rationale was developed by Sperling (1960); data from one experiment are shown in Figure 13-1. A four-by-three matrix of letters and numbers was presented to observers for 50 milliseconds (0.05 sec). An example of such a matrix is the following:

7	I	V	F
X	L	5	3
B	4	W	7

Because observers could not be expected to report all twelve items, a tone of high or medium or low frequency served as an instruction to report the characters in the top or middle or bottom line of the matrix. The tone was sounded either before the display, at the moment the display ended, or at some time after the display (x -axis




Figure 13-1. Characters available from a three-by-four matrix of letters and numbers when the signal for a report occurs at various times relative to the display of the matrix. The signal was a tone of high, medium, or low frequency that instructed the observer to report the characters in the top, middle, or bottom row of the matrix. With increasing delays of the instruction tone, the accuracy of the observer's report approached the immediate-memory span for this task. The data are means across four observers. (Adapted from Sperling, 1960, Figure 7.)

in Figure 13-1). The mean number of characters correctly reported from the single line was multiplied by three, and taken as an estimate of the total reportable characters in the matrix (y-axis in Figure 13-1).

About 10 characters from the matrix could be reported when the instruction tone preceded the display or occurred as the display ended (delays of -0.10 and 0 sec); when the instruction tone followed the display with delays of 0.15 sec or longer, the proportion of reportable characters decreased with increasing delay. When the instruction tone was delayed by 1.0 sec, the number of reportable characters was about equal to the span of immediate memory in this task (this span, of about four to five items, is at the extreme low end of the usual five-to-nine item span of immediate memory presumably because of the relative complexity of the observer's task, including the mix of letters and numbers and the matrix format). The region between the data points and the immediate-memory span (shaded area in Figure 13-1) is assumed to represent the persistence of the sensory effects of the display; this is the extent to which the observer can still read particular rows of the matrix when given the appropriate instruction tone, even though the matrix is no longer present.

This experiment demonstrates that the sensory effects of a stimulus continue for a brief time after the stimulus has been presented. Similar effects have been demonstrated with auditory presentations of verbal stimuli, as in experiments in which different sequences of spoken letters were simultaneously presented in three auditory channels (left ear, right ear, and both ears: C. J. Darwin, Turvey, and Crowder, 1972). The effects of auditory stimuli appear to decrease more slowly, perhaps over several seconds, than those of visual stimuli. The auditory case has been called *echoic memory* (but should not be confused with the echoic class of verbal behavior in chapter 10; echoic memory involves the persisting effects of auditory stimuli, without regard to the nature of the listener's response to those stimuli, whereas echoic verbal behavior is defined by the correspondence between the auditory stimulus and the listener's vocal response). Obviously the finding that the sensory effects of stimuli continue for a brief time after the stimuli end has only a little to do with remembering over longer periods. Visual stimuli must be seen and auditory stimuli must be heard to be remembered, but remembering over minutes or days cannot be attributed to persisting sensory effects of stimuli.

Short-term memory: The role of rehearsal

The recall of a human learner immediately after the presentation of a verbal sequence, usually limited to about five to nine items, is the span of immediate memory (Jacobs, 1887; G. A. Miller, 1956). For example, one can probably correctly repeat the sequence 706294 after a single hearing, but a correct repetition of the longer sequence 541962836702 is unlikely. The immediate memory span has sometimes been described as the limit of what can be held in consciousness. What is tacit when we say that we are conscious of something at least includes responding under the control of our own behavior with respect to events. We report that we are seeing something or hearing something or remembering something by saying we are conscious of it.

The limit on the span of immediate memory provided the historical basis for studies of what has come to be called short-term memory or, in some usages, primary memory (J. Brown, 1958; L. R. Peterson and Peterson, 1959; Waugh and Norman, 1965; see also Daniels, 1895; W. G. Smith, 1895; Murray, 1976). A major question is whether the verbal response is directly under the control of prior stimuli or is instead under the indirect control of other intervening behavior, such as the behavior called rehearsal. One way to address this question is with a task that prevents the learner from engaging in such behavior between the presentation of some items and the opportunity for their recall. The tasks used to measure the span of immediate memory incidentally create such conditions. If a long sequence of items is presented, the learner cannot simultaneously rehearse the early items and listen to the later ones. Some of the later items therefore may interfere with or prevent the behavior upon which recall of earlier items depends.

Consider now the short-term memory experiment of L. R. Peterson and Peterson (1959; see also J. Brown, 1958). Vocal stimuli consisted of a sequence of

three consonants and then a three-digit number. The learner's instructions were to begin counting backwards by threes from the number and then, when a signal light flashed, to recall the consonants. A trial might start with the spoken items, CHJ 506; the learner then counted backwards, 506, 503, 500, 497, etc., until the signal light, which set the occasion for naming the consonants. The time from presentation of the consonants to onset of the signal varied from 3 to 18 seconds. Under these conditions, recall accuracy decreased with increasing delay until, at delays of 15 or 18 seconds, the proportion correct was less than 10%. In two other conditions, learners were allowed periods of vocal or silent rehearsal before they were given the number from which they were to count backwards. Figure 13-2 compares data from the original procedure (open circles) and data with 3 seconds of either silent rehearsal (filled triangles) or vocal rehearsal (filled squares) before the backwards-counting task. The way the learners responded to the consonants determined how well they were able to recall them. Recall is determined more by our behavior with respect to past stimuli than by the stimuli themselves.

These and related findings led to theories that proposed two distinct types of memory: short-term memory and long-term memory (e.g., Shiffrin and Atkinson, 1969). Items are said to be rapidly lost from short-term memory unless they are maintained there through rehearsal; by some means, perhaps through rehearsal itself, items in short-term memory are sometimes transferred to a more permanent long-term memory. Only a few items at a time can be held in short-term memory, but the capacity of long-term memory is virtually unlimited. Data from more




Figure 13-2. Proportion of three-consonant items correctly recalled as a function of the delay between the presentation of the item and the opportunity for its recall. In one condition (open circles), rehearsal was prevented by instructing the learner to count backwards by threes from a number that immediately followed the three-consonant stimulus item. In two other conditions, the learner was allowed three seconds of either silent rehearsal (filled triangles) or vocal rehearsal (filled squares) before a number that initiated the backwards-counting task was presented. (Adapted from L. R. Peterson and Peterson, 1959, Figure 3 and Table 1.)

detailed analyses of performances obtained with short-term memory procedures, however, complicated this view.

The data obtained without allowing rehearsal (Figure 13-2, opened circles) at first seem to represent the gradual fading or decay of the items to be remembered. But several lines of evidence suggest this is not the case. Consider first the effects of varying the number of letters to be remembered, as in Figure 13-3 (Melton, 1963). A stimulus item containing 1, 2, 3, 4 or 5 consonants and then a three-digit number were presented visually; the learner read the consonants aloud, and then began counting backwards by threes or fours until a visual signal set the occasion for recall. When the item consisted of only a single consonant, recall accuracy remained high even over a 32-second interval. The proportion of items recalled decreased more rapidly over time as the number of consonants increased. But which set of data might then represent the time course of short-term memory? If recall of a single consonant is nearly perfect, then it might simply be assumed that as other consonants are added each consonant interferes to some extent with recall of the others.

The two extreme possibilities in these procedures are perfect recall and no recall at all. Perhaps the backwards-counting task only partially prevents rehearsal, and the intermediate outcomes occur because the learner does have some opportunities for rehearsal even while engaging in that task. Because more rehearsal is presumably necessary for recall of more consonants, it follows that recall accuracy at a given delay should decrease with the number of consonants. Short-term memory thus seems to depend mainly on the learner's opportunity to respond to the stimulus items (cf. Crowder, 1976, p. 196). It may be relevant that the recall




Figure 13-3. Proportion of consonant items correctly recalled as a function of the delay between presentation and recall, with number of consonants as a parameter. The consonant items and a three-digit number were presented visually; the learner read the consonants aloud and then began counting backwards by threes or fours until the presentation of a visual signal for recall. (From Melton, 1963, Figure 2.)

probability of three consonants was greater in the Melton experiment (Figure 13-3) than in the Peterson and Peterson experiment (Figure 13-2): Melton's learners recited the consonants aloud and were therefore guaranteed at least one rehearsal before they began the backwards counting, whereas the Petersons' learners did not do so. One case used visual stimuli and the other used vocal stimuli, but the opportunity for immediate responding to the consonants was probably more important than the modality difference.

Once we recognize that each item may affect the recall of other items, we may ask how the numbers in the backwards-counting task come to affect recall. Do they prevent rehearsal, or do they interfere with recall in some other way? If each number produced by the learner during backwards counting is regarded as an item, does recall vary with time or with number of items? This problem was addressed by surveying memory experiments in which delays between item presentations and recall included various numbers or rates of intervening items (Waugh and Norman, 1965). Data from many studies indicated that the number of intervening items during the delay was more important than its duration. For example, learners were given a sequence of numbers at one per second or four per second, and then were asked to name the number at a particular location in the sequence. Recall probability after four intervening items at four per second was about equal to that after four intervening items at one per second, even though the delay between the item and the opportunity for recall was about one second in the former case and four seconds in the latter. On the other hand, recall probability was greater after a delay with two intervening items at one per second than after a delay with eight intervening items at four per second, even though both delays lasted about two seconds. If any rehearsal occurs while items are presented, items presented slowly should permit more rehearsal than items presented rapidly. Thus, how can the preceding outcome be reconciled with an account of short-term memory in terms of the effects of restricting rehearsal?

Still another problem is that recall probability changes over the first few trials of a short-term memory procedure (Keppel and Underwood, 1962; Kincaid and Wickens, 1970). For example, the recall probability of a three-consonant item on the first trial of a short-term memory session is nearly 1.0 at a delay of either 3 or 18 seconds; over the next three to six trials, the probability of recall decreases, but the decrease is much greater when the delay is 18 seconds than when it is 3 seconds. In other words, the relation between recall probability and the delay imposed by an interpolated task, as in Figures 13-2 and 13-3, is absent at the beginning of sessions and builds up over several trials. This is an example of proactive inhibition (cf. chapter 12): the learning of the consonants on the first trial interferes with the recall of other consonants on later trials. We will have more to say about such proactive effects later.

A change in stimulus items can temporarily eliminate the proactive effect in short-term memory experiments. For example, recall probability increases substantially for a trial or more after such variations in stimulus items as changing the items to be remembered from letters to numbers or from numbers to letters. Similar effects also occur with words as stimulus items (e.g., with a change in semantic

categories, as from furniture items to food items: Wickens, 1970). These procedures may be more relevant to defining verbal classes (cf. chapter 11) than to analyzing short-term memory.

Short-term memory procedures are usually designed to prevent rehearsal rather than to encourage it. It is therefore curious that rehearsal has played such a substantial role in interpreting the results of such procedures. We cannot specify a single short-term memory function, nor can we say confidently what the role of rehearsal is. Furthermore, demonstrating the effects of interference among items, both within and across trials, makes it inappropriate to speak of short-term memory as the passive decay of items. Nevertheless, the behavior of recalling an item after its uninterrupted rehearsal is different from recalling the item if there has been no opportunity for such rehearsal. This difference, together with the relatively limited span of immediate memory, justifies the distinction between short-term memory and other types of memory.

Long-term memory:

Interference and the nature of forgetting

Short-term memory is studied over periods of seconds or at most minutes, whereas long-term memory is usually studied over hours, days, or perhaps even years. It has been argued on the one hand that short-term and long-term memory are separate types of remembering with different properties, and on the other that short-term and long-term memory are merely extremes on a single continuum determined by variables that affect remembering (e.g., Melton, 1963; Tulving and Madigan, 1970). We will not attempt to resolve the controversy.

When the number of items in a list exceeds the span of immediate memory, a single presentation is not sufficient for the learner's recall, even if recall follows immediately after the presentation ends. If we wish to study the remembering of such lists, our procedures will necessarily include the repeated presentation of some or all of the items to be learned. The amount to be remembered may be a more important basis for distinguishing between short-term and long-term memory than the period of time over which remembering is measured. For our purposes, long-term memory will include those cases in which the items to be remembered exceed the immediate memory span. Such remembering cannot be based on uninterrupted rehearsal between presentation and recall, requires repeated presentations of the to-be-remembered items, and thus is likely to be studied over relatively longer time periods than those common to research on short-term memory.

Data from a classic study of long-term memory, by Ebbinghaus (1885), are shown in Figure 13-4. The remembering of nonsense-syllable lists was assessed by the method of savings (cf. chapter 12) at 20 minutes, 1 and 8.8 hours, and 1, 2, 6 and 31 days after original learning. The largest decrease in the proportion remembered (or, in other words, the most forgetting) occurred shortly after original learning. Even after 31 days, however, the savings on relearning exceeded 20%. Ebbinghaus entertained the alternative possibilities that memories deteriorated over




Figure 13-4. Ebbinghaus's forgetting curve. Ebbinghaus learned and then relearned 13-syllable lists. The time between original learning and relearning varied, and the saving upon relearning was taken as a measure of remembering (cf. chap. 12). Forgetting was substantial shortly after original learning (the first data point is at 20 minutes), but even after 31 days the savings were still greater than 20%. (Adapted from Ebbinghaus, 1885, chap. VII.)

these time periods or that they remained intact but were gradually overlaid by or hidden beneath other memories. (In a more contemporary vocabulary, the distinction is expressed as one between decreasing *availability* of items with time, as items are lost from memory, and decreasing *accessibility* with unchanged availability, as items remain in memory but become harder and harder to retrieve.)

These alternative views were the source of various theories of forgetting. Theories based upon the passive decay of memories, sometimes called *trace* theories, often assumed correspondences between remembering and hypothetical processes in the nervous system. Theories based upon competition among different memories, sometimes called *interference* theories, were more likely to rely upon environmental and behavioral variables. Nevertheless, the notion of memories as traces that faded or decayed over time was implicit in many early generalizations (e.g., Jost's Law: if two associations are of equal strength but different ages, the older one will lose strength or be forgotten more slowly than the newer one: Jost, 1897). The problem with such generalizations is that what is remembered may vary with the criterion used for learning and the way in which remembering is measured. For example, forgetting as measured by savings on relearning a list may differ from that measured by recall or recognition of items on the list. But an easily learned list is not necessarily better remembered than a list learned with difficulty. When the mastery of the items on two lists is equated by common learning criteria, the lists are forgotten at approximately equal rates, even if they differ in difficulty or the time to meet the learning criterion (Underwood, 1964).

The effects of sleep on memory provided one way to distinguish between accounts of memory in terms of decaying traces as opposed to interference. There is no convincing evidence for learning during sleep (e.g., Emmons and Simon, 1956), but less forgetting of what is learned while awake occurs during sleep than

during an equal waking period (J. G. Jenkins and Dallenbach, 1924). One interpretation was that events during waking were more likely than events during sleep to interfere with what had been learned. Other accounts, however, grew out of theories such as memory *consolidation*, which argued that memory was relatively impermanent immediately after learning, and that a period of time was required after learning for memory to become fixed or consolidated (e.g., Glickman, 1961). According to this theory, various events could disrupt the consolidation of memory, whereas sleep had no effect on or might even facilitate consolidation.

Consolidation theories were influenced by the phenomenon of *reminiscence* (Ballard, 1913; McGeoch, 1942; see also Kamin, 1957). Reminiscence, most likely to be observed with incompletely learned materials, is an increase in recall probability as time passes since the end of learning, later followed by the usual decrease in recall called forgetting. Reminiscence has been observed over minutes in some experiments and over successive daily sessions in others, and depends on the details of the learning and recall procedures. Perhaps the most important implication of reminiscence is that there is no such thing as a single function that describes remembering or forgetting; there even are circumstances in which recall becomes more rather than less likely with the passage of time. In any case, consolidation theories have been controversial (e.g., McGaugh and Dawson, 1971), and studies have usually been on effects of events such as electroconvulsive shock on animal learning rather than effects in human verbal learning (but see Tulving, 1969). Thus, accounts in terms of interference, in one form or another, have tended to dominate the analysis of long-term memory.

Initial attempts to treat forgetting in terms of interference assumed that events between learning and recall were the major source of interference. In other




Figure 13-5. Recall of the items on a list as a function of the number of lists previously learned. Each circle represents data from a different experimental study. Forgetting was greater as the number of previous lists increased; in other words, forgetting depended on proactive interference from the learning of the earlier lists. (Adapted from Underwood, 1957, Figure 3.)

words, the important events were assumed to act retroactively: events that preceded recall affected what had been learned at an earlier time. The critical finding for analyses in terms of interference was that earlier learning had substantial proactive effects on what was more recently learned, influencing the forgetting of more recently learned material (Underwood, 1957; Underwood and Postman, 1960). Figure 13-5 summarizes data from a variety of experiments, and shows how the forgetting of a list varies as a function of the number of lists learned earlier (Underwood, 1957).

The observation of proactive interference was significant because much data on human memory over more than half a century had been obtained from practiced learners who had served in experiments involving the learning of many lists. Ebbinghaus himself typified that circumstance, and even learners who did not participate in experiments with many different lists under many different conditions were ordinarily given practice lists before the experiment proper began. From those studies, the estimate of forgetting over 24 hours had been about 75%; Underwood's analysis showed that most of this forgetting is attributable to the proactive interference from the learning of earlier lists, and that without such interference forgetting is only about 25% (zero previous lists in Figure 13-5). In other words, a predominant cause of forgetting is that older learning interferes with remembering what is more recently learned.

This analysis showed that the remembering of a verbal list learned in the laboratory varied with the previous learning of other verbal lists in the laboratory. Once it was noted that a substantial amount of forgetting could be attributed to proactive interference from related material learned earlier in the laboratory, it was reasonable to ask how much of the forgetting that was left could be accounted for by proactive interference from sources outside the laboratory (Underwood and Postman, 1960). Even someone in an experiment on verbal learning for the first time enters with an extensive history of verbal behavior. If learning an earlier list in the laboratory can have some effect, then it would be surprising if verbal learning outside of the laboratory did not act in the same way. The nature of transfer is also relevant: depending on the relations among the different tasks, some kinds of previous learning might have larger or smaller proactive effects than others on the remembering of material learned after. The implications are tantalizing: perhaps there would be little forgetting or no forgetting at all if it were not for proactive interference (in fact, retroactive effects also contribute to forgetting).

One interpretation of proactive interference assumed that first-list items are in some way unlearned during second-list learning, but that, in a kind of spontaneous recovery, the first-list items gradually return to displace second-list items (cf. Barnes and Underwood, 1959). The plausibility of such accounts was reduced by comparisons of recall and recognition of first-list and second-list items at various times after learning, and by findings that proactive interference could vary with relations among individual items. Proactive interference is a phenomenon, but the details of its operation remain to be worked out. In any case, proactive interference does not explain forgetting; it is a name occasioned by the variables that affect forgetting (cf. Crowder, 1976, pp. 194-195).

*Overview: Metamemory as
differentiated and discriminated
remembering*

We have identified three memory classes: the persisting stimulus effects called iconic memory; the momentary recall, extendable through rehearsal but limited in capacity, called short-term memory; and the durable remembering called long-term memory. Our iconic memory of a visual item will depend on whether we have looked at it, and our short-term memory of the item will be influenced by how we have rehearsed it, and our long-term memory of the item will vary with the ways in which we have responded to the item. We need not quibble over whether remembering is behavior. What is more important is that, with the possible exception of iconic memory, remembering can be learned. Not only do we learn patterns of rehearsal that extend our short-term memory, and mnemonic techniques that influence our long-term memory, but we also learn to judge the properties of our own remembering in various circumstances (cf. Flavell, Friedrichs, and Hoyt, 1970). In other words, short-term and long-term memory are differentiated classes of remembering, and these classes are discriminated with respect to the circumstances in which they occur; how we remember will depend at least partly on past consequences of remembering. Metamemory might therefore be defined as the differentiation and discrimination of one's own remembering.

Consider the short-order cook working on two eggs over easy, two scrambled with bacon and one scrambled with ham, two orders of waffles, and a stack of pancakes. As each dish is handed out to be served and as each new order is called in, some items must be forgotten while others continue to be remembered. The earlier breakfast orders over the course of that morning, not to mention those on previous days, are potentially a substantial source of proactive interference. Nevertheless, this task can be and often is performed with considerable skill. The performance of such tasks, in which a set of items is continually updated by dropping some and adding others, is sometimes called *running memory* or *working memory*. Other examples are keeping track of the cards played in a card game and following the changing statistics during an athletic event (most appropriately, running memory of the races run, laps completed, relative positions of competitors, and other information at a track meet). In such circumstances, what can be forgotten may be as important as what is remembered.

Some studies have shown that learning some verbal items can be affected by instructing the learner to forget other items (e.g., Bjork, 1970). Interference by a given item is reduced by an instruction to forget that item. As might be expected, the reduction is greater when the instruction to forget precedes the presentation of the item than when it follows the item. But what sort of behavior occurs when the learner is instructed to forget? Are the items somehow deleted from memory, or do they remain while the learner comes to discriminate those to be remembered from those to be forgotten?

One experiment (Waugh, 1972) was based upon the effects of list length on recall probability of common one-syllable words. In free recall, the probability that

recall of the item is then called *retrieval* from storage. An item that has been stored is said to be *available*, but it is said to be *accessible* only if it can be retrieved.

These stages are analogous to storing index cards in and then retrieving them from a file, or storing information in and then retrieving it from the memory banks of a computer. The important difference is that we cannot examine any place where remembered items are retained in the way we can examine the contents of a file drawer or the fine structure of a computer core. Even what is remembered remains unspecified. We can examine an index card or the electrical charges in a computer, but we cannot find a remembered item within the learner's head. Although the language of storage and retrieval is an effective analogy for what happens in remembering, we must recognize its metaphorical status. In fact, one objective of some research on memory is to explore the limits of the metaphor, or the range of conditions over which it holds. An analysis of these limits may determine the circumstances in which we may appropriately tact the functional properties of remembering in the language of storage and retrieval.

By definition, a remembered item is one that has been stored and retained and retrieved. The failure to remember an item may occur because the item was not stored in the first place, or because the item was stored and then lost from storage during retention, or because the item was not retrievable at the opportunity for recall. In this section, we examine how remembering can be affected by events during these three phases.

Storage: Encoding and levels of processing

What the learner remembers is not so much the stimulus itself as the learner's response to it. The learner's response will inevitably differ from the stimulus, even when the two are in the same modality (as when the learner repeats a spoken verbal item). Remembering, in other words, is not simply reproduction of the stimulus. Even the immediate response to a stimulus cannot be interpreted in this way. The issue is long-standing. For example, the following by the Greek philosopher Theophrastus dates from about 300 B.C.:

...with regard to hearing, it is strange of him (Empedocles) to imagine that he has really explained how creatures hear, when he has ascribed the process to internal sounds and assumed that the ear produces a sound within, like a bell. By means of this internal sound we might hear sounds without, but how should we hear this internal sound itself? The old problem would still confront us. (Stratton, 1917, p. 85)

A more contemporary version of this point is the following:

Suppose someone were to coat the occipital lobes of the brain with a special photographic emulsion which, when developed, yielded a reasonable copy of a current visual stimulus. In many quarters this would be regarded as a

triumph in the physiology of vision. Yet nothing could be more disastrous, for we should have to start all over again and ask how the organism sees a picture in its occipital cortex. (Skinner, 1963; see also Skinner, 1976, p. 74)

As in the analysis of stimulus control, the problem of remembering is not to be solved by trying to follow the stimulus into the organism; rather, we must discover how to characterize the ways in which the organism behaves with respect to the stimulus.

The learner's behavior with respect to the stimulus to be remembered has been called *encoding* (e.g., Melton and Martin, 1972). Consider the following experiment (R. Conrad, 1964). In one part, the learner was to name spoken letters presented in background noise. In the second part, the learner was to write down in order a sequence of six letters that had been presented visually. In the first task, errors typically occurred along dimensions of common acoustic properties; for example, *V* was more likely to be confused with rhyming letters such as *B* or *C* than with other letters such as *N* or *X*. The same types of errors occurred in the second task; errors were again more likely along dimensions of common acoustic properties than of common visual properties (e.g., for straight lines versus curvature, written *V* has more in common than *N* than with *C*: cf. Gibson, 1965). Even though the second task used written stimuli and written responses, the learner is said to have encoded the stimuli acoustically and not visually. We need not assume that the learner's responses were subvocal speech; that remains a reasonable possibility, but more important is that the learner's remembering was based upon acoustic rather than visual dimensions of the stimulus letters.

The structure of memory depends on the dimensions along which stimuli are encoded. For example, some tasks may favor encoding based on semantic properties (defining technical terms), others may favor encoding based on formal properties such as modality (learning spellings or pronunciations), and still others may favor encoding based on temporal properties (following the score of a musical composition). The context in which a given item appears, the order or spacing or duration of items, the formal or semantic relations among items, and the instructions given to the learner all are properties that the learner might encode in some way. Furthermore, encoding may vary from time to time both within and across tasks, and may be based on combinations of properties as well as on single dimensions. It should not be surprising that we can identify different types and properties of encoding in different types of learning tasks. We will consider a few examples, in the context of the two major classes of encoding called *substitution* and *elaboration*.

Simple substitution is the most straightforward encoding, and corresponds to familiar examples of codes (e.g., as in Morse-code learning: cf. F. S. Keller, 1958). Consider learning a sequence of the binary digits, 0 and 1 (G. A. Miller, 1956). A sequence such as 010001101011101001 substantially exceeds the span of immediate memory. Each group of three binary digits, however, can be replaced by a single octal digit, according to the following list:

000 - 0	100 - 4
001 - 1	101 - 5
010 - 2	110 - 6
011 - 3	111 - 7

The binary sequence then can be coded as the octal sequence 215351, which is within the span of immediate memory. Encoding groups of items so that the number of items to be remembered is reduced has been referred to as *chunking*. The acoustic encoding of written letters can also be regarded as substitution, in that there is a unique correspondence between spoken and written letters.

A second major type of encoding is called *elaborative* encoding, as when a CVC nonsense syllable or a sequence of consonants is transformed into a word or a phrase. For example, a learner might rehearse the nonsense syllable BOH as "BOTH without a T," or the consonant sequence QBF as "Quick Brown Fox." Such verbal instances of elaborative encoding have been called natural-language mediation (Prytulak, 1971). Another variety of elaborative encoding is visual imagery, discussed in the context of mnemonics (cf. Paivio, 1969, 1971; Paivio and Madigan, 1968). Elaborative encoding does not have the unique correspondence between stimulus items and the learner's responses that characterizes substitution, and thus an encoded item may not always uniquely occasion the appropriate response. For example, the learner who encodes the word "baggage" according to the method of loci by imagining a tower of suitcases standing in front of the fireplace might well say "luggage" at recall.

Encoding, whether by substitution or elaboration, is almost inevitably selective. Depending on the task, the learner's responses are occasioned by some stimulus properties and not others. For example, for verbal stimuli presented visually, the learner's response may be unaffected by type face or by size; similarly, for stimuli presented vocally, the response may be unaffected by the speaker's voice or by loudness. The encoding called reading is more likely than other kinds of encoding based on the physical properties of written words. It is not easy to look at a word without reading it. We may call some of our responses to verbal stimuli semantic encoding. For example, the responding occasioned by the word "log" is not the same in the context of "log book" as in the context of "log cabin" or "log jam" (cf. chapter 11). This encoding may affect retrieval. If "jam" appears in one list in the context of "log jam," in later lists it will be more likely to be recognized in some contexts, such as "traffic jam," than in others, such as "raspberry jam" or "jam session."

These categories of encoding are neither exhaustive nor mutually exclusive. The treatment of mnemonics, for example, combined substitution, in a number-to-consonant code, and elaboration, in visualizing objects at successive locations according to the method of loci. We should not expect an exhaustive listing, because types of encoding are as unlimited as the different ways in which humans can respond to events in the world. For the same reason, we should not expect any one type of encoding to be invariant across different learners or different tasks. Different learning histories and different contingencies upon remembering

guarantee variability in the ways that we each encode stimuli (cf. Anderson and Bower, 1973).

Different types of encoding differ not only in their likelihoods of occurrence in different circumstances, but also in the likelihood that what is encoded will be remembered. Some kinds of encoding seem more superficial and therefore less memorable than others; this dimension of encoding has been called *level of processing* or *depth of processing* (Craik and Lockhard, 1972; Glanzer and Koppenaal, 1977). Fewer words from a list are recalled after a task requiring responses to structural properties of words (e.g., crossing out vowels, copying the words, counting the number of letters) than after a task requiring semantic responses (e.g., assigning words to categories, discriminating among plant and animal names). The implication is that the greater the depth at which or to which an item is processed, the more likely it is to be remembered. We are certainly more likely to remember items to which we have responded in rich and novel and varied ways. The richness and novelty and variety can have many sources, of which semantic structure and visual imagery are important examples. As with so many concepts we have encountered, we must recognize level of processing as descriptive and not explanatory. It describes relations between types of encoding and the likelihood of remembering; it does not explain that remembering. In other words, we speak of semantic encoding as deeper processing simply because we are more likely to remember a semantically encoded item than a structurally encoded one. The issue of how we remember becomes one of what is remembered.

Retention: The question of memory reorganization

According to the metaphor of retrieval and storage, once an item has been encoded, it may be stored. But how does an item get stored, and what happens to the item after it has been stored? At this point, the metaphor of storage and retrieval does not help us much. Retention is implicit in the concept of storage, and yet we have already noted that we cannot expect to find the to-be-remembered item inside the learner. We have not even specified what an item is: a word? a sentence or logical proposition? an association? a semantic or syntactic structure tagged with various markers corresponding to its temporal and relational origins? It is only because the item emerges at recall that we assume that it has somehow existed throughout retention. That may be like assuming that a pianist at the keyboard releases Beethoven's sonatas from their storage in the piano, or, more simply, releases sounds that have been stored in the strings; whoever searches for the sonatas or the sounds by disassembling the piano will be disappointed.

We need to resolve this point, however, to ask whether what is remembered changes during retention. We mentioned earlier the theory of consolidation, which argues that what is learned becomes fixed or consolidated in memory over some time following learning. Another concept is that of incubation, which assumes that remembered events and relations are spontaneously reorganized over time (perhaps especially during sleep), sometimes in such a way that their combination constitutes




Figure 13-6. Examples of changes in remembered visual forms as a function of verbal labels. (Adapted from Carmichael, Hogan, and Walter, 1932, Figure 1.)

the solution to a long-standing problem (e.g., as in some examples of scientific creativity: Hadamard, 1949).

One type of study of memory change during retention used the reproduction of remembered material over time (e.g., the successive retelling of stories: Bartlett, 1932). An example is provided in Figure 13-6 (Carmichael, Hogan, and Walter, 1932). With a variety of visual figures, the remembered figure changed over retention as a function of the way in which it had been named. The most explicit theoretical case for these changes in memory during retention grew out of gestalt psychology (Köhler, 1929). The gestalt laws of perception (e.g., the law of closure, which stated that incomplete figures tend to be seen as completed) were assumed to operate on what is remembered as well as on what is seen. But changes during retention could not be distinguished from those that might have occurred during encoding and storage or at retrieval. If, because of verbal encoding, a visual stimulus were remembered as a word rather than a form, then we would expect the learner who remembered the word at recall to draw the corresponding picture. These examples suggest that there is no need to postulate any gradual reorganization of visual memory traces. We find, therefore, that we have little to say about what happens during retention. Perhaps that is why we usually speak of the metaphor simply as storage and retrieval instead of storage, retention and retrieval.

Retrieval: Cue-dependency and accessibility

We introduce retrieval with a quotation:

In order to understand retrieval processes, some basic principles must first be accepted. One of the most important of these was formulated by St. Augustine more than 1500 years ago; we cannot seek in our memory for anything of which we have no sort of recollection; by seeking something in our memory, "we declare, by that very act, that we have not altogether forgotten

it; we still hold of it, as it were, a part, and by this part, which we hold, we seek that which we do not hold." (Tulving and Madigan, 1970, p. 460)

In other words, we are not likely to be able to remember an item or event if we are not already able to discriminate, in some way, some properties of the item or event to be remembered. Sometimes these properties are specified by instructions, as when we are asked which of two items came earlier in a list, or whether we have met someone before, or where we were last year on the night of June 22. In these cases, the "parts we hold" are the items or a face or a date. At other times, circumstances define these properties, as when we cannot remember a name that we have to look up in the telephone book, or the article that we were supposed to buy at the store, or where we put our keys. Here, the "parts" are what we do remember of the person or place, or our presence in the store, or the activity that requires the keys.

An item that is stored is said to be *available* in memory. The item, however, may or may not be remembered. When it can be remembered, it is said to be *accessible*; when it cannot be remembered, it is said to be *inaccessible* (Tulving and Pearlstone, 1966). The accessibility of an item depends on the stimuli or cues present at the moment of recall. The influence of these stimuli on remembering is called *cue-dependency*. For example, visualizing a place that was used in the method of loci reinstates one condition that existed at the time an item was imagined at that place. Similarly, reciting the alphabet may help us to remember a forgotten name, to the extent that production of the person's initial may reinstate one condition for saying that name. The most important feature of remembering at recall may be the production of behavior similar to the learner's behavior during encoding and storage (e.g., Fisher and Craik, 1977).

Consider the experiment illustrated in Figure 13-7 (Tulving and Psotka, 1971). The recall of six 24-word lists was compared under three conditions. Each list was made up of sequences of four words in each of six semantic categories (e.g., military ranks: captain, corporal, sergeant, colonel; living places: hut, cottage, tent, hotel; earth formations: cliff, river, hill, volcano); the categories differed from list to list. One condition was recall of each list immediately after it had been presented (original learning). Another condition was recall after the presentation of all six lists (uncued recall). The third condition was recall after presentation of all six lists, but with the category names provided (cued recall). Although uncued recall showed the typical effect of interference from intervening lists, cued recall showed that the words had not been forgotten; when the category names were provided, recall was roughly equivalent to that after original learning. In other words, according to the metaphor of storage and retrieval, the words were available but inaccessible during uncued recall; the category names in cued recall made them accessible.

A similar cue-dependency was illustrated in chapter 12, in the treatment of conditions under which recall probability is greater than recognition probability. We may also note that the effective stimuli need not be those present during learning. For example, a rhyming word may be an effective cue for recall, even




Figure 13-7. Number of words recalled from six 24-word lists under three conditions. Each list was made up of groups of four words in six semantic categories (e.g., insects, metals, tools), and was presented visually three times in succession at a rate of one word per second. Categories differed from one list to the next. Data for original learning were obtained from recall of each list immediately after its three presentations. Data for noncued recall were obtained after all six lists had been presented; they are shown as a function of the number of lists that had intervened since the given list was presented. Data from cued recall were also obtained after all six lists had been presented, but the learners were given the names of the word-categories in each list. Cued recall was roughly equivalent to recall after original learning. Thus, during noncued recall the words had been available but not accessible. They were made accessible by the category names during cued recall. (Data from Tulving and Psotka, 1971, as presented in Tulving, 1974, Figure 2.)

though that word was not presented during learning and even though the learner's encoding of the original item had been primarily semantic. But recall is still likely to be maximal when encoding and recall cues are based on common properties of the to-be-remembered item (Fisher and Craik, 1977).

An implication of cue-dependency is that much forgetting can be attributed to differences between the conditions at encoding and storage and the conditions at recall. The presentation of some items from a list should therefore make conditions at recall more like conditions during learning. One paradox is that relative to recall of an entire list, the presentation of some of the items from the list lowers

rather than raises the probability that the learner will recall the remaining items (Slamecka, 1969). It is clear that, at least in some lists, relations among items have less effect on recall than other properties of the items.

Retrieval sometimes has been discussed in terms of the metaphor of search (Yntema and Trask, 1963). The learner is said to search through the memory store for items having particular characteristics until the appropriate item is found. The store might be analogous to a stack of index cards in a cabinet or to the memory banks of a computer. (The metaphors for human thinking have typically kept up with contemporary technologies; thus, in earlier times these metaphors were clock-word systems or hydraulic mechanisms: cf. McReynolds, 1971.) But the metaphor has problems in dealing with how rapidly and accurately humans can discriminate between knowing something and not knowing it (Kolers and Paley, 1976; see also Wason, 1965). Why should search ever take longer when an appropriate item exists to be found than when such an item does not exist (and therefore when the search must examine every item in storage)? Aside from the problem of how rapidly we can say that we do not know something, the metaphor also seems implausible because of the vast capacity of human memory; if nothing else, a search through memory storage at every instance of remembering seems inefficient. (Yet the increasing size of the memory store with experience, and thus the greater difficulty of finding any particular item, could be used as an argument for proactive interference; proactive interference is presumed to act more at retrieval than at storage.) Cue-dependency, however, suggests the possibility of a more restricted search, the limits of which are defined by the conditions at recall. The metaphor of search, though not likely to be adequate in general for accounts of human memory, may be relevant to some special classes of memory tasks.

In general, the implication of cue-dependency is that the more the conditions of recall are like the conditions of learning, the greater will be the likelihood of recall. These conditions include not only the to-be-remembered stimuli and the learner's behavior with respect to the stimuli, but also the situation within which the learning takes place. The special case of cue-dependency in which the relation between the learning and recall situations affects recall has been called *state-dependent learning*. The conditions studied include drug states and physiological states (e.g., electroconvulsive shock) as well as experimental settings (e.g., Overton, 1964; A. M. Schneider, 1967). For example, what was learned while intoxicated may be more likely to be remembered if the learner is again intoxicated than if the learner is sober.

Cue-dependencies and state-dependencies may have important practical implications. But we may not wish to advise the student who is preparing for an examination to study in the room in which the examination will be held. For one thing, what has been learned will be useless after the course is over if the student cannot remember it anywhere else. More important is that what is learned will be best remembered if many of the conditions for remembering arise out of the subject matter itself. A systematic subject matter is one in which each component prompts and is prompted by other components. To outline the subject matter,

therefore, is to create an encoding that will determine how well the various parts are remembered.

THE STRUCTURE OF MEMORY

We have considered various properties of memory. Many kinds of events and relations can be remembered. For example, the pigeon that pecks a key today because pecks yesterday occasionally produced food can be said to be remembering the past contingency between pecks and food. We have already questioned whether remembering of this sort is related in any simple way to remembering in human verbal learning. The vocabulary of memory may be occasioned by almost any situation in which current behavior is influenced by past events. We remember particular incidents, word definitions, contingencies, stimuli, language syntax, and our own behavior. Given this variety, we might expect properties of memory to vary with what is remembered. An example is motor memory (e.g., Baddeley, 1976, chapter 10). Two classes of motor memory have been distinguished: discrete skills, such as typing or changing gear in a manual-shift car; and continuous skills, such as tracking a continuously moving target or steering a car. The distinction is important mainly because discrete skills can be forgotten, whereas continuous skills rarely are (one does not forget how to swim or to ride a bicycle). These two types of skills seem closely related. Thus, if a memory difference can be observed even in this comparison, we must assume that other differences are likely over the range of different types of remembering. This section briefly surveys two classes of memory distinguished by what is remembered.

Episodic and semantic memory

When we speak of our remembering in everyday situations, we are usually concerned with particular incidents that occurred at particular times and places. This kind of remembering is called *episodic* memory, in that it involves the recall of episodes from our own life. Episodic memory can be contrasted with *semantic* memory, which is our remembering properties of language: word usages, idioms, and the functional aspects of grammar (Tulving, 1972; see also G. A. Miller, 1972, on lexical memory). We cannot understand a sentence if we do not remember what the words mean. Research on human verbal learning and memory seldom gives the learner completely new items. The words in a serial list are ones the learner has seen before, and even nonsense syllables are presented in the familiar alphabet. The learner's task is then to recall the items that appeared on a particular list in a particular setting. That list constituted an episode in the learner's life, and therefore the learner's recall can be characterized as episodic memory. The nature of such remembering may be different from that of remembering word meanings or mathematical relations, which are not recalled on the basis of a date or place of occurrence. (Such a difference is suggested by those cases of amnesia involving the forgetting of personal history without loss of language.)

The study of semantic memory usually deals with the structure of what is

remembered rather than with the functional properties of such remembering (cf. Anderson and Bower, 1973; Shimp, 1976). Some relevant research was presented in chapter 11, in the context of psycholinguistics. For example, a learner is more likely to remember semantic than syntactic properties of a sentence in a text (Sachs, 1967). Remembering a particular sentence that appeared in a particular text is episodic memory. But semantic structure is defined by how we generalize across syntactic transformations of sentences; one way we determine the nature of semantic and syntactic classes is by examining what is remembered. The correspondence between what is presented and what is remembered (recalled or recognized) defines structural classes in memory in the same way that the correspondence between the behavior that is reinforced and the behavior that is generated defines the structure of operant classes (cf. chapter 6). In other words, using remembering to study the structure of verbal behavior might be more profitable than using verbal behavior to study the nature of remembering.

A variety of experiments have examined the structure of what is remembered. Analyses of how stories and texts are remembered have led to accounts of the hierarchical structure of components of the story plot or text (e.g., Bower, 1976; McCoon, 1977). Analyses of the remembering of words and texts in different languages have shown that the structure of the semantic classes of bilingual speakers extends across language boundaries, as when the learner's remembering of meaning is independent of the language in which items are presented (e.g., Kollers, 1966; McLeod, 1976). Analyses of the implications that learners can derive from remembered and logically related sentences have shown how items of information can be organized into an integrated structure (e.g., Moeser, 1976; Moeser and Tarrant, 1977). Such studies inevitably combine episodic and semantic components of memory, because they include the remembering of both specific dated items and general structural relations, but they show also how an analysis of remembering depends on what is remembered.

SUMMARY

We treated the major memory classes called iconic and short-term and long-term, and the differentiated and discriminated remembering called metamemory. We examined the phases of an episode of remembering in terms of the metaphor of storage and retention and retrieval. We considered some properties of different types of remembering called episodic and semantic memory. It seems redundant to summarize these topics here in greater detail, because the essential vocabulary of memory has been included in the section headings of this chapter. These headings can be listed and remembered in a variety of ways: by rote rehearsal, by mnemonic techniques, or by reviewing the relations among particular experiments and concepts. Undoubtedly the reader who learns this material merely as a sequence of words will not remember it as well as the reader who has responded in other ways; that difference has been spoken of as depth of processing. We remember contingencies, words and texts, specific events, structures, and our own behavior. We have not explained memory, but we have seen that what is remembered is essential to the definition of important functional units and dimensions of behavior.

Cognition and Problem Solving

COGNITIVE PROCESSES

Visual imagery

Processing stages

Mental representations

PROBLEM-SOLVING

FOURTEEN

The Indo-European root *gno-* is the source through Old English of *know* and *knowledge* and through Germanic of *cunning* and *can*, in the sense of being able to. Its Latin derivative is *gnoscere*, to know or to get acquainted with, and this is a root of *cognition*, *ignorant*, *recognize* and *notice*. Synonyms of *knowledge* are often related to sensory language. For example, the Indo-European *weid-*, to see, leads to *view* and *vision* through the Latin *videre*, to *idea* through the Greek *eidos*, and to *guide* and *wisdom* through various Old English and Germanic forms. The word *see* itself implies understanding, as in the phrase *I see*, and has such relatives as the word *insight*.

Early in this text, before embarking on our treatment of the effects of stimuli and contingencies, we distinguished between structural and functional problems in the analysis of behavior (chapter 1). Structural problems are those concerned with properties of response and stimulus classes; these classes are the fundamental units of behavior. We have now considered a variety of response classes and stimulus classes, both verbal and nonverbal, some defined by contingencies (reinforcers, operants, and discriminated operants) and others defined by what is occasioned or what is remembered (semantic and syntactic classes). Structural problems are complemented by functional ones; in fact, function enters into the definition of structure (as when we say that several responses are members of the same class because the responses enter into similar functional relations). We have considered such functional relations as those between responses and their consequences in reinforcement and punishment, or between discriminative stimuli and their occasioned responses in stimulus control, or between the presentation and the later remembering of an item in memory. We can define functional relations only in terms of the behavioral units that enter into them, and therefore the definitions of structure and function are reciprocal. With some arbitrary structure as a starting point, we determine functional relations, and then use these relations to refine the

definition of the structural units, then make those units in turn the basis for a more rigorous description of function, and so on.

We indicated, in our initial discussion of function and structure, that concern with functional problems tends to be correlated with a behavioral vocabulary, whereas concern with structural problems tends to be correlated with a cognitive vocabulary. In this text, we began with an emphasis on the experiments and vocabulary that had grown out of the behavioral tradition. The behavioral treatment was extended to verbal behavior, which began the transition to topics that are the primary concern of contemporary cognitive psychology. Those topics were examined in the context of psycholinguistics, verbal learning, and memory. This chapter further illustrates research methods and issues in cognitive psychology; it then proceeds to a brief treatment of problem-solving, in which structural analyses converge with the functional properties of stimulus control and contingencies.

COGNITIVE PROCESSES

We already have spoken of some classes of behavior called cognitive processes (e.g., imagining or visualizing). Cognitive processes are private events: we cannot see what someone else is thinking or imagining. But imagining, like walking or talking, is something we do. Some might argue that the different dimensions of these private events do not justify calling them behavior. Yet our treatment of behavior has not been limited to particular movements; operants, for example, are not defined by response topography. Certainly it is difficult to say just what a person does when imagining (cf. chapter 10 on private events). Nevertheless, it is plausible to assume that such behavior shares something with the behavior of looking at things in the environment (we can discriminate our imagining from our seeing; when we fail to do so, we are said to have hallucinations; cf. Skinner, 1953; Paivio, 1975). The difficulty is that this behavior is relatively inaccessible to anyone but the one who engages in it. We may recall the recommendation to convert psychological nouns to verbs (Woodworth, 1921): instead of "cognition and thought," we should say "knowing and thinking" (see also Malcolm, 1971).

Consider the private event called paying attention or attending to something. In chapter 7, we treated attention as the control over responding by some dimensions of a stimulus but not others. That treatment established a word from the everyday vocabulary as a technical term, but the term remains in the everyday language as a tact of something we do. Attending is a kind of response. The distinction might be maintained in the technical vocabulary by restricting the noun, *attention*, to dimensional stimulus control, and reserving the verb, *to attend*, for analyses of the response class. The technical usage, however, is not our present concern. Instead, the issue is what kind of a response attending is: how can we identify its occurrence, and how can we measure its properties?

It is not enough simply to ask whether someone is attending. The problems of introspection, in which individuals report private events, have been amply

demonstrated in the history of psychology. Introspection was not reliable. Different individuals, and the same individual at different times, reported events in different ways. Attending presented particular difficulties: how does one attend to one's attending? The difficulties were inevitable, because the language of private events must be based at some point upon what is publicly accessible to the verbal community (cf. chapter 10). A private event is available only to the person behaving. This does not mean that private events are outside the boundaries of a behavioral analysis; it does imply, however, that we learn the language of private events on the basis of consistencies in the relations between private events and public events. Imagining an object, for example, must have something in common with seeing the object.

In any case, the response of attending is not a movement. For visual stimuli, it may superficially seem analogous to looking toward or even pointing at something. But if attending were simply equivalent to eye movements, situations in which we look without seeing would be excluded. We do sometimes look without seeing, however, as when we are daydreaming or "lost in thought." We sometimes then say that we were not paying attention. The behavior called attending may also share something with what we do when we say we are watching for or looking for something. Such behavior has been called perceptual set or readiness (e.g., Bruner, 1957). For example, what we notice or attend to if we are scanning a page to find a name will not be the same as if we are scanning to find a definition.

We have already seen that when a response class is not measurable by movement, we may be able to measure other properties of the behavior, such as its duration or its latency (cf. chapter 11). This rationale was extended to the response of attending by Sperling and Reeves (1976). Observers were instructed to look at a fixation point in the visual field. To the left of this point, letters were presented in rapid sequence one at a time. To the right, a sequence of numerals was similarly presented. The instructions to the observer were to attend to the letters until seeing a particular letter (e.g., *B*), and then to attend to the numerals and report the first numeral seen. Humans can perform this task without eye movements and with rates of stimulus projection in excess of twenty per second. Because the temporal relation between the critical letter and each numeral was known, the time taken to switch from attending to letters to attending to numerals could be derived from the number reported. Observers were unable to report the number that appeared simultaneously with the critical letter; instead, they reported a number that appeared some fraction of a second later. These durations provided a distribution of latencies or reaction times for a change in attending. Like the reaction times of more obvious responses, such as button presses, these reaction times increased with task difficulty.

Visual imagery

Even if attending is a response, what about more complex private events such as visualizing or imagining? We considered imagery in connection with mnemonic techniques in chapter 13. The methods for measuring imagery are necessarily also

indirect, but varied experimental techniques for the study of imagery have been explored (Paivio, 1971). Here again, temporal measures have been particularly effective. For example, observers were shown pictures of pairs of three-dimensional figures, with instructions to report whether or not the figures were the same (Shepard and Metzler, 1971). When the figures were the same, one was rotated relative to the other, as illustrated in Figure 14-1. For such figures, the latency of the report that they were the same was linearly related to the difference in orientation, in degrees of rotation. In other words, mental rotation, or rotation of an image, has some of the same properties as the rotation of an actual object: in both cases, the time taken to perform the rotation is proportional to the distance through which the object is to be rotated. But by treating the image as a thing we must not be misled into using it to explain behavior; the imagining of the object, like the rotation itself, is something the observer does. The temporal properties of rotating real and imagined objects are similar because of the correspondence between what the observer does when seeing an object and what the observer does when imagining it. Imagining is visual behavior in the absence of the visual stimulus.

It has been suggested that Watson, the founder of behaviorism, may have denied the existence of images because he himself was incapable of visual imagery (Skinner, 1959). The recognition that such a denial was irrelevant to a behavioral analysis was slow in coming. Whether the image has some physical status inside the observer's head is beside the point. The advantage of saying that we imagine an object over saying that we have an image of the object is simply in its emphasis on




Figure 14-1. Pairs of figures from a study of mental rotation. Observers were instructed to report whether or not the two figures were the same. In A, the two figures are the same, but the right figure is rotated 80° from the left one in the picture plane. In B, the two figures are different, and cannot be matched by any rotation. (From Shepard and Metzler, 1971, Figure 1.)

what we do. What is at issue is not the status of images, but rather the role of such private events (or cognitive processes) in accounts of behavior. Our primary concern must be to determine their own properties rather than to base accounts of other kinds of behavior upon them. A behavioral account need not deny the existence of such events, but it limits the conditions under which such events can properly serve as explanations of behavior (cf. Paivio, 1975, p. 287).

The example of mental rotation used visual stimuli. The response of visualizing also occurs in the absence of such stimuli, as we saw in the case of visual encoding as a mnemonic technique. In such cases, there is even more temptation to speak of the image as a thing seen rather than of imagining as behavior. In the "photographic" memory of someone with eidetic imagery, for example, reports of the details of past scenes are accompanied by eye movements similar to those when scanning a current visual scene (e.g., Haber, 1969). Yet to say that the reports or the eye movements are caused by the remembered image does not explain anything; an account of imagining must be derived from the behavior of seeing and not from the stimulus properties of a private event.

Consider some additional research on visual imagery. In one experiment (Moyer, 1973), observers were given pairs of words (e.g., *mouse* and *elephant*), and were instructed to choose the larger. The greater the difference in size between the two named classes, the shorter were the reaction times (e.g., the response to *truck-cat* was faster than that to *book-chair*). The words themselves did not differ systematically in size, and the relation between size difference and reaction time was similar to that with pictures rather than words as stimuli. These and related experiments (e.g., Paivio, 1975) suggest that this task involves imagery occasioned by words; such behavior is presumably analogous to the visual encoding discussed in chapter 13.

In another type of experiment, learners are given pairs of statements such as "*A* is taller than *B*" and "*A* is shorter than *C*," or "*X* is to the left of *Y*" and "*Z* is to the right of *Y*," and then are instructed to indicate whether such statements as "*B* is taller than *C*" or "*Y* is to the left of *Z*" are true or false (e.g., Brooks, 1968; Huttenlocher, 1968; Scholz and Potts, 1974; see also Huttenlocher and Strauss, 1968; Moeser and Tarrant, 1977). Latency measures and the differential effects of concurrent visual and nonvisual tasks suggest that such performances are better described as the visualizing of spatial relations than as verbal responding occasioned by the statements. The private construction of a spatial array in such a task seems to parallel the behavior of moving objects into various spatial arrangements. These experiments, therefore, identify visualizing as a behavioral class and then demonstrate that this class is a component of a particular performance.

Treating imagery as a kind of behavior as opposed to something that an observer "has" or "does not have" raises the possibility that visualizing can be taught. An artist presumably learns some of this behavior in progressing from sketching a live model to sketching without a model. Shaping of visualization might proceed by gradually dimming a scene as an observer describes it and by gradually

increasing the time between the presentation of the scene and the observer's description. Systematic studies of such phenomena, however, are not available in the research literature.

Processing stages

As we have seen, temporal measures have been an important feature of the analysis of cognitive processes. Besides showing that private events take time, some cognitive analyses have been concerned with demonstrating how tasks can be decomposed into their separate components or stages.

One of the oldest ideas in experimental psychology is that the time between stimulus and response is occupied by a train of processes or *stages*—some being mental operations—which are so arranged that one process does not begin until the preceding one has ended. This *stage theory* implies that the reaction-time (RT) is a *sum*, composed of the durations of the stages in the series, and suggests that if one could determine the component times that add together to make up the RT, one might then be able to answer interesting questions about mental operations to which they correspond. (Sternberg, 1969, p. 421)

Early treatments attempted to identify particular stages, such as a discriminative stage (e.g., by subtracting the simple reaction time to a single visual stimulus from the discriminated reaction time to one of two visual stimuli). At various times, such stages as stimulus registration, decision, and response initiation were postulated (cf. Sternberg, 1967). Two problems, however, complicated stage theories. One was that stage durations determined in different laboratories and even at different times in the same laboratory were highly variable, presumably because most tasks can be performed in a variety of ways. Second, and perhaps more important, stage durations were calculated on the assumption that the presence or absence of one stage had no effect on the duration of other stages. This assumption might not be justified. For example, consider one task made up of stages *A-B-D-E* and another made up of stages *A-B-C-D-E*. If the duration of *B* or *D* in one of these contexts differs from that in the other context, the duration of *C* will not be given by the difference between the latencies on the two tasks.

One solution is to devise a task in which one of the stages is repeated a specified number of times (e.g., *A-B-C-C-C-C-D-E*). In this instance, the immediate contexts of stages *B* and *D* are not altered by changing the number of repetitions of stage *C*. Thus, the increase in latency produced by each added repetition of *C* may be taken as equivalent to the duration of *C*. This was the rationale for experiments designed to explore memory search or memory scanning as a possible stage in remembering (Sternberg, 1966, 1969). An observer was first given a set of digits (e.g., 1, 3, 4, 9), called the positive set. Then, when digits from 0 through 9 were

presented visually, the observer's instructions were to press one button if the digit was a member of the positive set and another if it was not. The latency from digit presentation to button press was measured over successive conditions in which the number of digits in the positive set was varied.

How might this task be performed? To speak of search is in this case a metaphor, but what properties should such a metaphorical search have? The observer's response is determined by whether the presented digit matches one of the digits in the positive set. One stage of this task, therefore, might consist of comparisons between the presented and the positive-set digits. Let us begin by assuming that the comparisons are made one at a time instead of all at once, or, in other words, that the observer engages in *serial search* instead of *parallel search*. Serial searches may vary in at least two different ways. First, the observer might search the positive-set digits either in a varying and random order or in a fixed order (e.g., beginning to end); we will speak of *random-order* and *fixed-order* search. Second, the observer may discontinue comparisons as soon as a match is found, or may complete all comparisons before attending to matches; the former is called *self-terminating* search and the latter *exhaustive* search. The various types of search have different quantitative implications, as illustrated in Figure 14-2.




Figure 14-2. Quantitative implications of exhaustive and self-terminating patterns of serial search. In exhaustive search, reaction time as a function of list length has the same slope for positive as for negative items; in self-terminating search, the slope for positive items is half that for negative items. In exhaustive search, reaction time for positive items is independent of the positions of items in the list; in self-terminating search, reaction time as a function of position in the list depends on whether search is in random or fixed order. (From Sternberg, 1969, Figure 4.)

Consider first the relation between reaction time and list length. In exhaustive search the presented digit, whether in the positive set or not, is compared with all the positive-set digits. For example, if each comparison takes 50 milliseconds, each item added to the positive-set digits should add this duration to both positive and negative reaction times (negative reaction times are those to digits not in the positive set); the slopes of these two functions in this case will be equal (the two functions will thus be parallel, but they might differ by a constant if positive responses take more or less time than negative responses). With self-terminating search, however, all comparisons will be made for negative digits, because no matches will be found, whereas for positive digits a match will be found, on the average, after only half of the comparisons. In this case each digit added to the positive set will add 50 milliseconds to negative reaction times but only 25 milliseconds on the average to positive reaction times; the slope for positive items will therefore be half that for negative items.

Next consider the relation between reaction times to a positive item and the position of that item in the positive set. If search is exhaustive, all comparisons will be made, and position in the list will have no effect (in Figure 14-2, horizontal lines for list lengths of 3 and 7 items). If the order of comparisons is random, self-terminating search will also produce this independence between reaction time and the position of the positive item in the list. But if the order is fixed, from the beginning to the end of the list, reaction times in self-terminating search will vary with position; matches to items early in the list will occur more quickly than those to items late in the list.

Figure 14-3 presents some data from this type of procedure. Positive and negative reaction times, averaged across eight observers, are shown as a function of number of digits in the positive set. Reaction times increased roughly linearly, at about 38 milliseconds per item, for both positive and negative items. These reaction times are consistent with the metaphor of search, and imply that search in this task is exhaustive rather than self-terminating. The times per item, however, are too short to justify the assumption that the comparisons involve the observer's subvocal rehearsal of the items in the positive set.

It may seem paradoxical that search should be exhaustive rather than self-terminating. If a match has occurred, why should the observer continue with comparisons through the end of the list? One possibility suggested by Sternberg is that the comparisons and the determination of a match are themselves separate stages in this task; if switching between them also takes time, then exhaustive comparisons may be more efficient with relatively short lists than a self-terminating succession of the individual comparisons each followed by the determination of a match.

This task is high specific, and even changes in detail change the characteristics of the data (cf. Baddeley, 1976; Crowder, 1976). More complex variations of these procedures, for example, produced reaction-time functions characteristic of self-terminating rather than exhaustive search; in addition, the functions relating




Figure 14-3. Mean reaction times for positive and negative items as a function of the number of items in the positive set, for eight observers. Errors (usually within 1 or 2% of total responses) were excluded from the data. Reaction time (RT) as a function of number of items (N) is roughly described by the equation: $RT = 400 + 38N$, where RT is in milliseconds. (From Sternberg, 1966, Figure 1.)

reaction times to position of the item in the positive list were consistent with random-order search for some observers and fixed-order search for others (Sternberg, 1969; cf. Figure 14-2). The metaphor of search is strengthened by the quantitative detail of the data, but its range of application is limited (e.g., different effects may be obtained if the number of items in the positive set exceeds the span of immediate memory; for this reason, the metaphor has not been extended to memory in general: cf. chapter 13). Furthermore, the metaphor is based on the assumption of serial search rather than parallel search, in which items are compared simultaneously instead of successively. Parallel search need not be restricted to cases in which the search takes a fixed amount of time regardless of the number of items in the positive set; comparisons might become slower as the number of items increased. Thus, it is possible to develop an account in terms of parallel search that is equivalent in outcome to those outlined for serial search (e.g., Townsend, 1971).

Despite these limitations, the advantage of this type of approach is in its close adherence to the quantitative details of performance. An adequate description of the private events called processing stages will necessarily depend on the combination of studies that examine particular stages in isolation (as in the example of

attending with which this section began) and those concerned with the dynamics of more complex tasks (as in this instance).

Mental representations

In their emphasis on the temporal properties of attending and visualizing and searching, these treatments primarily have dealt with defining the structure of private events. They have been less concerned with functional issues such as the conditions that occasion these events or the consequences they may have. But, as we saw earlier, there are other aspects of structure, such as sequential and hierarchical organization (e.g., consider the sequential and hierarchical components of completing a college major: the major consists of individual courses the order of which may be constrained by prerequisites; the courses may be decomposed into smaller units consisting of assignments and examinations; these in turn may be further reduced to reading particular passages or answering certain questions, again perhaps in specified orders; and so on). Treatments of concept learning and strategies in problem-solving sometimes proceed by showing how correspondences may develop between the structure of the environment and the structure of behavior (e.g., Garner, 1974; Haygood and Bourne, 1965; Restle, 1962).

The development of cognitive structure has been a central theme in the work of Piaget (e.g., Piaget and Inhelder, 1969; see also Fischer, 1975; Ginsburg and Oppen, 1969). We can provide here only a brief and inevitably oversimplified discussion. The several phases of child development in Piaget's system (sensory-motor, preoperational, concrete operational, and formal operational) correspond to a progression from relatively simple relations between motor responses and their consequences (as in reaching for and manipulating objects) to complex relations that depend both on correspondences between verbal and nonverbal behavior and on correspondences between environmental structure and behavioral structure.

Piaget speaks of the development of such correspondences in terms of *accommodation* and *assimilation*: the child must accommodate to the constraints imposed upon her or his behavior by environmental structures and contingencies, but these structures and contingencies are assimilated to that extent that they become incorporated into the child's behavior. What is assimilated, however, is sometimes said to be a structure, such as a mental representation. Cognitive development, therefore, is assumed to be the gradual enrichment of representations of the world; these representations then become the basis for behavior. Bruner (1964) has elaborated on this view by distinguishing among three varieties of representation: *enactive* representation, corresponding to motor performance (as in thinking about riding a bicycle or walking from one place to another); *iconic* representation, corresponding to imagining in the several sensory modes (as in visualizing a scene or recalling a melody); and *symbolic* representation, corresponding to verbal behavior (as in formulating a description of some event or deriving a mathematical proof).

The phenomenon called *conservation* illustrates some aspects of cognitive structure dealt with in terms of representations. If the water in a short wide beaker is poured into a tall narrow beaker, a preschool child is likely to report that there is more water in the second beaker than there was in the first. This report is based on the higher level at which the water stands in the tall narrow beaker than in the short wide one. In Piaget's account of development, this type of response is characteristic of the preoperational phase; one property of this phase is that responding is based upon single properties of environmental events rather than upon relations among properties. Later, in the concrete operational phase, the child says that the amount of water is the same in one beaker as in the other. Such reports are restricted to specific instances, and it is only in the subsequent formal operational phase that the child begins to speak abstractly of these relations as the conservation of matter.

One way to deal with this phenomenon is to attribute the increasing sophistication of the child's performance to successive refinements in the child's mental representations of events in the world. As the child moves from actually manipulating objects to imagining or talking about or thinking about manipulating those objects, the structure of the child's mental representations are said more and more closely to approximate the structure of the contingencies that operate for manipulating objects in the world. But the structure of these mental representations was derived from the correspondence between behavior and environmental contingencies. It would therefore be inappropriate to assume that the behavior is explained by these representations. The representations do not exist in the child in any but a metaphorical sense; instead, they exist in our own discriminative behavior as we observe the child, and correspond to the consistencies and correspondences that we observe in her or his behavior.

The development of conservation in children is in any case not incompatible with properties of behavior that we have already considered. Several progressions must occur in parallel, and each is part of a hierarchy in which more complex classes of behavior are based upon simpler classes. At first, the child's discriminations are based upon a single dimension of objects (in this instance, height); discriminations based upon two or three dimensions (area and volume) come later. At the same time, the child's vocabulary of tacts becomes more finely differentiated (from *big-little* to *bigger-smaller* and then to *taller-shorter*, *fuller*, and so on). Tacts under the control of particular instances become the foundations for the more general classes of tacts we call abstractions. And so on. Saying whether or not the quantity of water in one container equals that in another is not the same as choosing the contents of one or the other container or filling two containers to equal volumes, and none of these is equivalent to a statement of the principle of conservation of matter. On the one hand, it is important to recognize how very different these classes of behavior are; on the other, it is important to recognize that consistencies across these different classes are inevitable consequences of the way in which verbal and nonverbal behavior is related to events in the world.

PROBLEM-SOLVING

We have examined only a few samples of research in cognition. Cognition is, after all, not learning; yet many of the issues in cognition are closely related to problems in learning. As the last experimental topic in this text, we turn briefly to problem-solving. Problem-solving has sometimes been treated in the context of cognitive approaches, with emphasis on the structure of problems. An analysis of problem-solving also raises functional questions, as when we are concerned with conditions that may make the solution of a problem more or less likely.

Let us begin with a structural approach to problem-solving in the Hobbits-and-Orcs problem (Thomas, 1974; see also Greeno, 1974). Three hobbits and three orcs are traveling together. The orcs will not leave the hobbits behind, but they will overpower them if ever any hobbits are outnumbered by orcs. The group reaches a river that must be crossed, and finds a boat that can hold only one or two creatures at a time. How do the hobbits organize the crossing so that no hobbits are ever outnumbered by orcs? An experiment explored the effects of feedback and other variables when humans were given this problem to solve.

This problem allows only a small number of possible moves, and its solution can economically be presented as in Figure 14-4. It is curious that, except for an alternative first move not shown and the two cases of branching moves at the beginning and end of the sequence, the only allowable alternatives to correct moves are moving backward through the sequence (this property of the problem was seldom recognized by solvers). For our purposes, it is sufficient to note that errors were more likely at some points in the sequence (states 321 and 110) than at others. These and other features of the data indicated that the problem-solver's performance did not consist simply of the separate moves. Instead, the solution was based on larger units consisting of sequences of moves that led to some intermediate arrangement of hobbits and orcs on the two sides of the river. Thus, this analysis is concerned with the structure of the problem and of the problem-solver's solution.

In what way, however, can the structure of the solution to one particular problem contribute to the solution of other problems? This question is necessarily relevant to those researchers concerned with teaching problem-solving skills within well-defined areas such as mathematics (e.g., Wertheimer, 1959). Analyses must proceed from single problems to transfer from one problem to another. Experiments on problem-solving, therefore, have often dealt with the effects of the problem-solver's history on problem solution.

An example is the Luchins' water-jar problem (Luchins, 1942; see also Luchins and Luchins, 1950). Given a water supply and three jars of different capacity, the problem is to produce a specified quantity of water. For example, if jars *A*, *B*, and *C* have respective capacities of 21, 127, and 3, how does one arrive at a quantity of water of exactly 100? In this instance, the solution is to fill *B* and then to pour out enough to fill *A* once and *C* twice. One group was given several successive problems that could be solved in this way, while a second group was not




Figure 14-4. Successive states in the solution of the Hobbits-and-Orcs problem. Three hobbits and three orcs must cross a river. They have a single boat that can hold only one or two creatures, and under no circumstances may any hobbits be outnumbered by orcs. Each state is represented by three digits: number of hobbits on starting side, number of orcs on starting side, and boat on starting side. In each state, the positions of the hobbits (H), the orcs (O), and the boat are shown below the digit code, and the creatures transported by the boat are shown in the transitions between states. Except for a transition to state 320 from the first state (crossing of one orc and the boat, not shown), no other moves in which one or more hobbits are not outnumbered by orcs are possible. (Adapted from Thomas, 1974, Figure 1.)

given these problems. Later, both groups were tested on problems such as the following: how does one arrive at a quantity of 20 if jars *A*, *B*, and *C* have respective capacities of 23, 49, and 3. The first group continued to solve the problem in

the same way, whereas the second group simply filled *A* and then poured out enough to fill *C*.

We might call this an example of negative transfer or problem-solving rigidity, but it may be more important that this is obviously a case in which current behavior is determined by the consequences of past behavior. The literature on problem-solving includes many variations on such situations (e.g., Birch and Rabinowitz, 1951; Duncker, 1945; Saugstad and Raaheim, 1960; Weisberg and Suls, 1973). For example, problem-solvers who had been instructed to mount a candle on a wall were given either a box of candles and a box of matches and a box of tacks, or separate stacks of candles, matches, tacks, and boxes (Duncker, 1945). The solution was to tack a box to the wall so that it provided a flat surface on which a candle could be mounted. Those who received the candles and matches and tacks in boxes, however, were much slower to solve the problem than those who received the boxes separately. Instances of failure to solve a problem that requires an unusual use of common tools or materials are sometimes described as cases of *functional fixity*; the commonplace items occasion responses appropriate to their everyday functions instead of the novel responses appropriate to the problem solution.

These experiments show how problem-solving can be affected by discriminative stimuli and by contingencies. But what then constitutes problem-solving? Consider an example from Skinner:

You have been asked to pick up a friend's suitcase from an airport baggage claim. You have never seen the suitcase or heard it described; you have only a ticket with a number from which a match is to be found among the numbers on a collection of suitcases. To simplify the problem, let us say you find yourself alone before a large rotary display. A hundred suitcases move past you in a great ring. They are moving too fast to be inspected in order. You are committed to selecting suitcases essentially at random, checking one number at a time. How are you to find the suitcase? You may, of course, simply keep sampling. You will almost certainly check the same suitcase more than once, but eventually the matching ticket will turn up. . . . A much more effective strategy is to mark each case as it is checked—say, with a piece of chalk. No bag is then inspected twice, and the number of bags remaining to be examined is reduced as rapidly as possible. Simple as it seems, this method of solving the problem has some remarkable features. . . . It is the use of the chalk which introduces something new. Marking each suitcase as it is checked. . . is constructing a discriminative stimulus. (Skinner, 1966, pp. 136-37)

Problem-solving, in other words, is behavior. The discriminative features of the situation define the problem, and the reinforcer is the solution of the problem. Often we solve problems by manipulating stimuli: looking something up, converting a verbal problem into a mathematical equation, and so on. It is important to recognize how we can affect our own behavior by modifying the stimuli in our environments. We still have much to learn about such processes (e.g., Catania, 1975; Rachlin, 1974). But now that we have seen how problem solution can be affected

by discriminative stimuli and contingencies, we have come nearly full circle. The experimenter interested in problem-solving might start simply by watching the behavior of the problem-solver. Like Köhler in his study of insight in apes, the experimenter will discover what behavior the problem-solver is already capable of. That is where our systematic behavioral treatment began, with the simplest behavioral operation: observation. We may now leave the experimenter to studies of stimulus presentations and contingencies and stimulus control.

Conclusion

STRUCTURE AND FUNCTION
LEARNING AND EVOLUTION

FIFTEEN

The word *psychology* comes from the Greek *psukhein*, to breathe, and *psukhe*, breath. Thus, when the psyche was said to leave the body of a mortally wounded warrior in the Homeric epics, the word may be understood to mean the warrior's breath and not his soul. A curious parallel exists between this word and *spirit*, from the Latin *spirare*, another word meaning to breathe. There is no evidence that these Greek and Latin words are etymologically linked. A similar relation between air and spirit exists in the Latin *anima*, originally a breath of air but later soul or spirit in such words as *animate* and *animosity*, and in the Greek *atmos*, vapor or air as it enters into the English *atmosphere* but breath or soul in the related Sanskrit *atman*.

We have considered the varied phenomena of learning. We began with experimental operations. The first and simplest was the presentation of stimuli. Stimuli produce responses, and the effects of stimuli on the probability of responses can be modified by the temporal patterning of stimulus presentations. We also saw that the stimulus-presentation operation could have other effects: stimuli can affect the way in which responses are ordered in time, in the classes of behavior called interim and adjunctive, and the elicitation of responses may affect their subsequent emission, in the phenomenon historically called exercise.

Just as responses may follow stimuli, stimuli may follow responses. We therefore dealt also with the effects of contingencies or the consequences of responding, in treating the consequential operations of reinforcement and punishment. The effects of these operations led us to distinctions among positive reinforcement, positive punishment, negative reinforcement (or escape and avoidance) and negative punishment (or omission procedures). In exploring these procedures, we noted the importance of some consequences, less dramatic than such biologically significant stimuli as food and water, that affected an organism's orientation in its environment. In this context, we examined sensory-motor learning, visual fixation, latent learning, and other phenomena not usually classified according to the effects of reinforcers or punishers.

We distinguished between reinforcement and punishment by whether response consequences raised or lowered the probability of responding, but we also recognized that distinguishing between the positive and negative cases of these operations was often arbitrary. As in reinforcement by a change in temperature, we sometimes could not consistently treat the positive and negative cases in terms of whether response consequences involved presenting or removing stimuli (e.g., presenting heat or removing cold). Instead, we treated these cases in terms of temporal relations between reinforced or punished responses and the other responses generated by these consequential operations. We noted that stimuli could have multiple effects, and that a stimulus that served as a reinforcer for a response might also have eliciting effects on that response or other responses. Thus, the task of an experimental analysis is to separate such different stimulus effects and to observe how elicitation and consequences can combine to determine the characteristics of a performance.

We also turned to the signalling effects of stimuli, in the operation of stimulus control. This operation can be superimposed either on the consequential operations of reinforcement and punishment, when we speak of discrimination learning, or on the operation of stimulus presentation, when we speak of respondent conditioning. Within these contexts, we explored a variety of behavioral phenomena and procedures: attention, stimulus-control gradients, reinforcement schedules, sensory preconditioning, autoshaping, and conditioned suppression, to mention only a few.

These operations seem to exhaust the possible behavioral relations that can be established among stimuli and responses. Our concerns therefore shifted to a different problem: characterizing the dimensions along which stimuli and responses vary. In discussing the processes of differentiation and discrimination, we saw that the relevant dimensions included not only relatively simple ones such as topography or intensity, but also complex ones involving the structure of stimuli and responses and the relations among them. We found that the correspondence between the classes of responses that had particular consequences and the classes of responses generated by those consequences was critical to defining behavioral classes, and we distinguished between classes defined by response properties, simply called operants, and classes defined by stimulus properties, called discriminated operants. In our examinations of concept formation and verbal behavior, we were forced to conclude that behavioral relations were primary, and that we could not appeal simply to physical measures to define such classes. There are responses, like attending, remembering, imagining and thinking, that are not easily observed directly, but we decided early that it would be inappropriate to define behavior solely in terms of movement.

Much of our account classified learning phenomena according to experimental operations. But there are no guarantees that imposing a particular procedure on a given organism's behavior will be effective. A procedure that leads one organism to learn may not so lead other organisms, and an organism led to learn under one procedure may not be led to do so under other procedures. Such findings do not invalidate our classifications, because the classifications are merely ways of naming phenomena and relating them to each other. That would not be so, however, if we

assumed that all learning was based on some one process or some small number of processes that act across all organisms and procedures.

The psychology of learning evolved through the examination of the different experimental outcomes of various behavioral procedures. As each procedure was considered in its historical turn, it was assigned an importance in proportion to its demonstrated effectiveness. At different times, the available experimental findings led to theoretical formulations dominated by laws of association or contiguity, rules of respondent conditioning, or principles of reinforcement and punishment. Sometimes this domination was so substantial that one or another process was presumed to be the fundamental and exclusive basis of all learning. Such formulations were inevitably open to challenge and controversy, because the phenomena of learning could not be accounted for exhaustively by any single process. As learning theories became more precise, they were necessarily accompanied by refinements and qualifications restricting the range of phenomena to which the theories could be applied. Recognizing the boundary conditions for learning was implicit in these restrictions.

Limits on learning pose no problems when they can be easily traced to an organism's sensory or motor capacities. For example, we are not surprised if certain stimuli are more likely to produce responses in some species than in others. We know that the visual and auditory systems of pigeons and bats make pigeons capable of visual discriminations impossible for bats and make bats capable of auditory discriminations impossible for pigeons. We are also not troubled by the different capacities for responding. Pigeons and bats fly differently because the anatomical structures of their wings are different. Although the examples are obvious, they are not trivial. They illustrate how much we take the different sensory and motor capacities of different species for granted.

Like sensory and motor capacities, the effectiveness of different stimuli as reinforcers or punishers varies across species. Once we recognized the relativity of reinforcement, it was no longer adequate simply to identify which reinforcers or which punishers might be effective for a given species. These consequences can only be defined in relation to the responses that produce them. Within a single species, a reinforcer effective for one response may not be effective for others. As demonstrated by phenomena such as food aversion and acknowledged by the concept of preparedness, we cannot specify the limits on learning in terms of stimuli alone or responses alone. Instead, we must express these limits in terms of the kinds of relations that can be established between stimuli and responses in a given species. Organisms may be predisposed to learn different relations among stimuli and responses in different situations.

The boundaries within which various experimental operations are effective in changing behavior do not alter the properties that define behavioral processes. It does not follow that we should begin to call a phenomenon by a different name if it turns out not to be as general as we once believed. If we argued that the sole principle of learning was demonstrated by a particular behavioral procedure or process, such as elicitation or reinforcement or stimulus control, then its generality would be challenged by the discovery of conditions that limited its applicability.

Once we recognize that all of the various operations and processes may be important in their different ways, the discovery of the limitations of each becomes an integral part of our analysis rather than a challenge to it.

STRUCTURE AND FUNCTION

Early in our account, we distinguished between structural and functional analyses (chapter 1; Titchener, 1898; Catania, 1972, 1973c). In a structural analysis, the relations among stimuli and responses are held constant while critical properties of one or the other are altered. As an example, we considered how textual units could be varied in teaching reading, and we noted as major concerns the hierarchical organization in the structure of text (as in relations among letters, syllables, words, phrases and sentences) and the correspondence between stimulus structure and response structure (as in relations between text and speech). In a functional analysis, the stimuli and responses of interest are held constant while the relations among them are altered; the interactions between behavior and environment are studied in terms of contingency relations among discriminative stimuli, responses and consequences. As an example, we considered how the effects of reinforcing correct responses or punishing errors might be compared in the teaching of reading.

Structural and functional problems are often interrelated. If we were designing an illustrated reader for use in an elementary school classroom, we would want to consider both the structural features of the reader, in the details of text and pictures, and its functional features, in ways in which the pictures might be related to the sentences that they illustrate. For example, a problem of stimulus control might exist if the pictures set the occasion for verbal descriptions that corresponded so closely to the accompanying text that a teacher could not tell whether a child was responding to the picture, reading the text, or doing both in some combination. This problem could be addressed by changing the relation between the pictures and the child's verbal responses (e.g., by making the pictures reinforcing consequences of rather than discriminative stimuli for reading). A systematic analysis might show that pictures, as consequences, maintain the picking up of books and the turning of pages, but that as discriminative stimuli they compete with the text in setting the occasion for particular verbal responses. Such an analysis might lead to a rationale for sequencing the relations among words and pictures in children's readers to maximize both stimulus control by the words and reinforcing effects of the pictures. At the same time, a structural analysis of the text and pictures could provide the basis for organizing the textual material, deciding on appropriate levels of complexity, and perhaps even determining the rate at which the proportion of text relative to pictures should increase (Harzem, Lee, and Miles, 1976).

Structure and function are not mutually exclusive; they are complementary, and there is no reason why structural concerns should not enter into a functional experiment, or vice versa (the analogous structure-function distinction in biology, between anatomy and physiology, is also sometimes ambiguous). However difficult it may sometimes be to draw the line between the two types of cases, the structure-

function dimension remains a basis for distinguishing among different research problems.

We indicated in chapter 1 that in the evolution of psychology the distinction between structural and functional approaches became correlated with and eventually confused with another distinction, that between the two psychological languages of cognitive psychology and behaviorism. The cognitivist prefers to summarize the organization of the organism's environment and its behavior in terms of structures that the organism knows, for which the language of mind and the mental is convenient. The behaviorist prefers to anchor accounts of action to the details of functional relations among observable events, to which the language of stimuli and responses is appropriate. Yet the relations between stimuli and responses are implicit in the cognitive vocabulary, because these relations are summarized in the correspondence between environmental structure and behavioral structure that is intrinsic in structural assumptions. Similarly, structure is implicit in the behavioral vocabulary, because it is summarized in the properties that functionally define the response classes called operants and discriminated operants. (Biology did not have an equivalent problem: the debate over mechanism versus vitalism, perhaps analogous to that between behaviorism and mentalism, remained reasonably independent of the distinction between anatomy and physiology: Hein, 1972; Catania, 1978.) In the development of psychology, the respective correlations of cognitive and behavioral positions with structural and functional problems made it difficult to recognize that the problems were different and therefore that these research areas might be complementary rather than mutually exclusive.

When the biologists of an earlier time debated structure and function, their concerns were not so much with whether structural or anatomical problems could be distinguished from functional or physiological problems as they were with whether one or the other problem should be given priority. The arguments were based on assumptions about how, in the evolution of a species, the functions of an organ might determine its structure or how, in the development of an organism, the structure of an organ might determine its functions. The problem was resolved not with the domination of one or the other position, but rather with the recognition that structure and function are mutually determined by environmental selection pressures: the structural and functional properties of organisms within a species are each to be understood in terms of the evolution of that species.

Analogous controversies over the primacy of structure or function exist in the history of the psychology of learning, although the issues have not ordinarily been expressed this way. For example, consider latent learning (chapter 4): a hungry rat explores a maze with an empty goalbox, and when it is later given food in the goalbox, it demonstrates that it has learned the maze by negotiating it as rapidly and with as few blind-alley entries as a rat that has always found food in the goalbox. The rat's running is functional; it occurs because of its relation to the food in the goalbox. The running itself is not learning; it is a performance that presumably reflects what has been learned. It was once important to distinguish

between learning and performance, but this distinction merely served as a basis for asserting the primacy of structure over function. Yet what is structure in this situation? The performance is a particular sequence of turns, and is structured in this sense. To say that a structure was learned, however, is not to say that the structure was a cause of learning. A theorist who wanted to reinstate the primacy of function over structure might then argue that the structure is learned because of the contingencies: particular turns at particular choice-points led to particular new locations in the maze, and the rat learned these functional relations even without food in the maze. Simply in negotiating the environment, the consequences of behavior matter. Then comes the argument that if contingencies define what is learned, they too cannot be a cause of learning. And so it goes. The theorist arguing for the primacy of structure was probably a cognitivist, and the theorist arguing for the primacy of function was probably a behaviorist, and their different languages were not likely to help matters.

Yet these structures and functions of behavior were both outcomes of learning. The problems can be resolved only by recognizing that behavioral structure and behavioral function are mutually determined by the relations between behavior and the environment. Both the structure and the function of behavior are to be understood in terms of their origins, and neither has primacy over the other.

LEARNING AND EVOLUTION

In the analysis of behavior, it has been argued that the problems of learning parallel those of evolution, because the selection or survival of patterns of behavior in an organism's lifetime has much in common with the selection or survival of individuals in the evolution of a species (e.g., Skinner, 1966b, 1975; Staddon and Simmelhag, 1971; Gilbert, 1972; Segal, 1972; Catania, 1978). Both cases must involve some kind of variation that provides the source materials on which evolution acts, and both must have some basis for selecting what survives. In each case, we must deal not with particular instances but rather with populations or classes of events. We speak of populations of organisms as species, and we speak of behavioral populations as response classes, such as operants, discriminated operants, and respondents.

It might be argued that our notions of response classes are much vaguer than the classes Darwin spoke of as species in his account of evolution (Darwin, 1859). Yet even though the word *species* is in Darwin's famous title, *On the Origin of Species*, Darwin knew that the term could not be given an unambiguous definition. In his book, he commented often on the problem of definition:

I look at the term species, as one arbitrarily given for the sake of convenience to a set of individuals closely resembling each other (p. 52); . . . the amount of difference necessary to give two forms the rank of species is quite indefinite (p. 59); [and] . . . we shall have to treat species in the same manner as those naturalists treat genera, who admit that genera are merely artificial

combinations made for convenience. This may not be a cheering prospect; but we shall at least be freed from the vain search for the undiscovered and undiscoverable essence of the term species. (p. 485)

In Darwin's account of evolution, the relations among populations of organisms could not be expressed adequately in terms of similar topographies (for example, males and females within some species differ more from each other in form than some pairs of organisms within completely unrelated species; the social insects in particular are striking instances).

For Darwin, the important basis for distinguishing among populations was descent. We define the relations among populations by looking at where they came from. Darwin's achievement, in fact, was more description than explanation. His account of evolution did not depend on any theory specifying the mechanics of evolution (his work predated genetics, and he even argued at times that acquired traits might be the source of variation on which evolution acted); rather, he had described the properties of evolution (Gould, 1975, p. 824). Darwin himself was skeptical about explanation: "It is so easy . . . to think that we give an explanation when we only restate a fact" (Darwin, 1859, pp. 481-482).

What has this to do with learning? The analogies between behavior and biology suggest that some solutions appropriate to biology will be appropriate to an analysis of behavior. We have seen the importance of descriptions of what happens in learning, and have given such descriptions priority over theories or mechanisms or models. For example, we found that the term *reinforcement*, which once served explanatory functions, now functions simply as a name for a particular behavioral phenomenon. Thus, we no longer ask whether reinforcement explains behavior. Instead, empirical questions about the generality of reinforcement are best regarded as questions about the range of circumstances to which this name can be appropriately applied. Similarly, the study of memory does not explain what is remembered; instead, what is remembered is the basis for determining behavioral structure, as in the analysis of the psychological reality of syntax and semantics. The theories and models will come and go, but the descriptions of behavior will remain (and it is fitting that we have sometimes found it appropriate to consider the results even of experiments conducted in an earlier century).

Darwin's treatment rejected the concept of immutable species in favor of classes defined by their descent or evolution. The psychology of learning has sometimes moved in the opposite direction. It sought to *explain* learning, by investing sources for responses (neural traces, associations, or cognitive structures). But that was getting it backwards. We should use the development of behavior to define behavioral classes. We should define behavioral classes in terms of their descent: where they came from or how they were learned.

To some extent, we already do that, though seldom explicitly. We distinguish innate behavior from behavior acquired through experience. We speak of behavior produced by stimuli as elicited and behavior occurring independently of stimuli as emitted. We say that responses engendered by stimulus presentations are operant respondent behavior and responses engendered by consequences are operant

behavior. We sometimes even offer an account of the origins of responses as an explanation of behavioral phenomena (e.g., as when behavioral contrast is analyzed as a combination of operant and elicited responses: chapter 8). But we were not satisfied; we did not recognize how much had been accomplished.

In the analysis of behavior, we deal with populations of responses. These populations are sometimes not well defined. The problem is not different from Darwin's. Darwin clearly recognized the arbitrary nature of the concept of species, but unambiguous definitions of species were no more critical to his account than unambiguous definitions of stimulus classes or response classes need be to a behavioral account. When we distinguish among words by the circumstances in which they are uttered (e.g., "Fire" as a mand, a tact, an echoic, or a textual response: chapter 10), we are simply distinguishing among classes of verbal responses on the basis of their origins. Ambiguous cases will necessarily occur because, just as organisms have many ancestors, responses have many origins. But there would be little need for an analysis if that were not the case.

One conclusion to which this account leads is that behavior is best represented not by all-or-none categories but rather by various dimensions along which processes can be located. Reinforcement and punishment are extremes on a continuum of contingencies that range from those that increase responding, through those that do not affect responding, to those that decrease responding. Positive and negative reinforcement are extremes on a continuum of relations involving the relative temporal locations of the stimuli present when responses are reinforced. Differentiation and discrimination are extremes on continua that represent how much response properties and stimulus properties contribute to the criteria of differential reinforcement. Operant discrimination and respondent conditioning are extremes on a continuum that relate stimulus-control operations to contingencies, including both those in which responses change stimulus probability and those in which responses have no effect on stimulus probability. Tacts and mands and other verbal classes are extremes on continua that define the relative contributions of antecedent stimuli and consequent stimuli to particular instances of verbal behavior. Syntax and semantics are extremes on a continuum that relates structural properties of verbal behavior to complex features of the verbal and nonverbal environment. Verbal recall and verbal recognition are extremes on a continuum of different sorts of activity entering into the classes of behavior that we call remembering. Behavioral and cognitive processes are extremes on a continuum that represents the relative accessibility or inaccessibility of things that organisms do.

The term *learning* has receded into the background in all of this; perhaps it has outlived its usefulness. We can modify behavioral hierarchies, shape new responses, establish discriminations, form concepts, and solve problems. Our understanding of these phenomena depends at least in part on how consistently we can speak of them. We have often recognized the ambiguities of the present vocabulary of behavior, and we can assume that this vocabulary will change with research progress. Yet in emphasizing behavioral operations and processes, it at

least adheres closely to what is done and what is observed in experiments on behavior. Whatever else is incorporated into any future psychology of learning, these operations and processes must provide its foundations.

References

[Bracketed numbers indicate the pages on which each reference is cited.]

- Adams, J. A. Psychomotor performance as a function of intertrial rest interval. *Journal of Experimental Psychology*, 1954, 48, 131-133. [283]
- Alleman, H. D., and Platt, J. R. Differential reinforcement of interresponse times with controlled probability of reinforcement per response. *Learning and Motivation*, 1973, 4, 40-73. [175]
- The American Heritage dictionary of the English language. Boston: Houghton Mifflin, 1969. [etymologies]
- Amsel, A. The role of frustrative-nonreward in noncontinuous reward situations. *Psychological Bulletin*, 1958, 55, 102-119. [188]
- Anderson, J. R. Verbatim and propositional representation of sentences in immediate and long-term memory. *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 149-162. [269]
- Anderson, J. R., and Bower, G. H. Recognition and retrieval processes in free recall. *Psychological Review*, 1972, 79, 97-123. [297]
- Anderson, J. R., and Bower, G. H. *Human associative memory*. Washington, D.C.: Winston, 1973. [325, 329, 335]
- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *Journal of Experimental Psychology*, 1956, 52, 145-161. [175]
- Anger, D. The role of temporal discriminations in the reinforcement of Sidman avoidance behavior. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 477-506. [111]
- Anisfeld, M., and Knapp, M. Association, synonymity, and directionality in false recognition. *Journal of Experimental Psychology*, 1968, 77, 171-179. [296]
- Antonitis, J. J. Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, 1951, 42, 273-281. [121]
- Appleyard, D. Styles and methods of structuring a city. *Environment and Behavior*, 1970, 2, 100-118. [146]
- Aslin, R. N. Development of binocular fixation in human infants. *Journal of Experimental Child Psychology*, 1977, 23, 133-150. [91]

- Atkinson, R. C. Optimizing the learning of a second-language vocabulary. *Journal of Experimental Psychology*, 1972, 96, 124-129. [288]
- Autor, S. M. *The strength of conditioned reinforcers as a function of frequency and probability of reinforcement*. Doctoral dissertation, Harvard University, 1960. In D. P. Hendry (Ed.) *Conditioned reinforcement*. Homewood, Ill.: Dorsey, 1969. [196]
- Azrin, N. H. Some effects of two intermittent schedules of immediate and non-immediate punishment. *Journal of Psychology*, 1956, 42, 3-21. [98, 185]
- Azrin, N. H., and Hake, D. F. Positive conditioned suppression: conditioned suppression using positive reinforcers as the unconditioned stimuli. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 167-173. [217]
- Azrin, N. H., and Holz, W. C. Punishment. In W. K. Honig (Ed.) *Operant behavior: Areas of research and application*. New York: Appleton-Century-Crofts, 1966. [97, 100]
- Azrin, N. H., Holz, W., Ulrich, R., and Goldiamond, I. The control of the content of conversation through reinforcement. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 25-30. [244]
- Azrin, N. H., Hutchinson, R. R., and Hake, D. F. Extinction-induced aggression. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 191-204. [83]
- Azrin, N. H., Hutchinson, R. R., and Hake, D. F. Attack, avoidance, and escape reactions to aversive shock. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 131-148. [57]
- Azrin, N. H., Hutchinson, R. R., and McLaughlin, R. The opportunity for aggression as an operant reinforcer during aversive stimulation. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 171-180. [83]
- Baddeley, A. D. *The psychology of memory*. New York: Basic Books, 1976. [334, 344]
- Badia, P., Suter, S., and Lewis, P. Rat vocalization to shock with and without a CS. *Psychonomic Science*, 1966, 4, 117-118. [52]
- Baer, D. M., Peterson, R. F., and Sherman, J. A. The development of imitation by reinforcing behavioral similarity to a model. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 405-416. [164]
- Ballard, P. B. Obliviscence and reminiscence. *British Journal of Psychology, Monograph Supplements*, 1913, 1 (No. 2). [322]
- Barnes, J. M., and Underwood, B. J. "Fate" of first list associates in transfer theory. *Journal of Experimental Psychology*, 1959, 58, 97-105. [289, 299, 302, 323]
- Barrett, J. E. Conjunctive schedules of reinforcement: II. response requirements and stimulus effects. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 23-31. [185]
- Barrett, J. E. Conjunctive schedules of reinforcement: III. a fixed-interval adjusting fixed-ratio schedule. *Journal of the Experimental Analysis of Behavior*, 1976, 25, 157-164. [185]
- Bartlett, F. C. *Remembering: a study in experimental and social psychology*. Cambridge: Cambridge University Press, 1932. [269, 330]
- Beale, G. H. Adaptation in paramecia. In R. Davis and E. F. Gale (Eds.) *Adaptation in microorganisms. 3rd Symposium of the Society for General Microbiology*. Cambridge: Cambridge University Press, 1953. [5]
- Bechterev, V. M. *General principles of human reflexology*. (tr. by E. Murphy and

- W. Murphy, from the fourth Russian edition, 1928) London: Jarrolds, 1933. [205, 207]
- Bem, D. J. Self perception: an alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, 1967, 74, 183-200. [216]
- Beninger, R. J., and Kendall, S. B. Behavioral contrast in rats with different reinforcers and different response topographies. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 267-280. [188]
- Berko, J. The child's learning of English morphology. *Word*, 1958, 14, 150-177. [274-275]
- Berlyne, D. E. *Conflict, arousal, and curiosity*. New York: McGraw-Hill, 1960. [89]
- Bernbach, H. A. Stimulus learning and recognition in paired-associate learning. *Journal of Experimental Psychology*, 1967, 75, 513-519. [296]
- Berryman, R., and Nevin, J. A. Interlocking schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 213-223. [185]
- Best, J. B. The photosensitization of paramecia aurelia by temperature shock. *Journal of Experimental Zoology*, 1954, 126, 87-99. [5]
- Bever, T. G., Lackner, J. R., and Kirk, R. The underlying structures of sentences are the primary units of immediate speech processing. *Perception and Psychophysics*, 1969, 5, 225-234. [261]
- Bewley, W. L. The functional stimulus in serial learning. In R. F. Thompson and J. F. Voss (Eds.) *Topics in learning and performance*. New York: Academic Press, 1972. [283]
- Bigelow, G. Fixed-ratio reinforcement of spaced responding. *Journal of the Experimental Analysis of Behavior*, 1971, 16, 23-30. [183]
- Bilbrey, J., and Winokur, S. Controls for and constraints on auto-shaping. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 323-332. [214]
- Bilodeau, E. A. (Ed.) *Principles of skill acquisition*. New York: Academic Press, 1969. [132]
- Bindra, D. Components of general activity and the analysis of behavior. *Psychological Review*, 1961, 68, 205-215. [18]
- Birch, H. G., and Rabinowitz, H. S. The negative effect of previous experience on productive thinking. *Journal of Experimental Psychology*, 1951, 41, 121-125. [350]
- Bjork, R. A. Positive forgetting: the non-interference of items intentionally forgotten. *Journal of Verbal Learning and Verbal Behavior*, 1970, 9, 225-268. [324]
- Blackman, D. E. Conditioned suppression or facilitation as a function of the behavioral baseline. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 53-61. (a) [218-219]
- Blackman, D. E. Response rate, reinforcement frequency, and conditioned suppression. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 503-516. (b) [218, 220]
- Blackman, D. E. Conditioned suppression and the effects of classical conditioning on operant behavior. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N.J.: Prentice-Hall, 1977. [215, 217]
- Blodgett, H. C. The effect of the introduction of reward upon the maze performance of rats. *University of California Publications in Psychology*, 1929, 4, 113-134. [86]
- Blough, D. S. New test for tranquilizers. *Science*, 1958, 127, 586-587. [119]

- Blough, D. S. Delayed matching in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1959, 2, 151-160. [309]
- Blough, D. S. Steady-state data and a quantitative model of operant generalization. *Journal of Experimental Psychology: Animal Behavior Processes*, 1975, 104, 3-21. [159]
- Boakes, R. A. Response decrements produced by extinction and by response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 293-302. [85]
- Bolles, R. C. Species-specific defense reactions and avoidance learning. *Psychological Review*, 1970, 77, 32-48. [110, 206]
- Bolles, R. C. *Theory of motivation*. (Second edition) New York: Harper and Row, 1975. [61]
- Boren, J. J., and Devine, D. D. The repeated acquisition of behavioral chains. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 651-660. [164]
- Bousfield, A. K., and Bousfield, W. A. Measurement of clustering and of sequential constancies in repeated free recall. *Psychological Reports*, 1966, 19, 935-942. [292]
- Bousfield, W. A. The occurrence of clustering in the recall of randomly arranged associates. *Journal of General Psychology*, 1953, 49, 229-240. [292]
- Bower, G. H. Analysis of a mnemonic device. *American Scientist*, 1970, 58, 496-510. [310, 312]
- Bower, G. H. Experiments on story understanding and recall. *Quarterly Journal of Experimental Psychology*, 1976, 28, 511-534. [335]
- Bower, G. H., McLean, J., and Meacham, J. Value of knowing when reinforcement is due. *Journal of Comparative and Physiological Psychology*, 1966, 62, 184-192. [190]
- Braine, M. D. S. On learning the grammatical order of words. *Psychological Review*, 1963, 70, 323-348. [275]
- Brandauer, C. *The effects of uniform probabilities of reinforcement on the response rate of the pigeon*. Unpublished doctoral dissertation, Columbia University, 1958. [172]
- Bransford, J. D., Barclay, J. R., and Franks, J. J. Sentence memory: a constructive versus interpretive approach. *Cognitive Psychology*, 1972, 3, 193-209. [269]
- Bransford, J. D., and Franks, J. J. The abstraction of linguistic ideas. *Cognitive Psychology*, 1971, 2, 331-350. [269]
- Breland, K., and Breland, M. The misbehavior of organisms. *American Psychologist*, 1961, 16, 681-684. [137]
- Brogden, W. J. Sensory preconditioning. *Journal of Experimental Psychology*, 1939, 25, 323-332. [208]
- Brooks, L. R. Spatial and verbal components of the act of recall. *Canadian Journal of Psychology*, 1968, 22, 349-368. [292, 341]
- Brown, J. Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology*, 1958, 10, 12-21. [316]
- Brown, P. L., and Jenkins, H. M. Auto-shaping of the pigeon's key-peck. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 1-8. [213]
- Brown, R. *A first language: the early stages*. Cambridge, Mass.: Harvard University Press, 1973. [274]
- Brown, R., and McNeill, D. The "Tip of the Tongue" phenomenon. *Journal of Verbal Learning and Verbal Behavior*, 1966, 5, 325-337. [239, 325]

- Bruce, R. W. Conditions of transfer of training. *Journal of Experimental Psychology*, 1933, 16, 343-361. [299]
- Bruner, J. S. On perceptual readiness. *Psychological Review*, 1957, 64, 123-152. [339]
- Bruner, J. S. The course of cognitive growth. *American Psychologist*, 1964, 19, 1-15. [346]
- Butler, R. A. The effect of deprivation of visual incentives on visual exploration motivation in monkeys. *Journal of Comparative and Physiological Psychology*, 1957, 50, 177-179. [89]
- Bykov, K. M., *The cerebral cortex and the internal organs*. (tr. by W. H. Gantt) New York: Chemical Publishing, 1957. [206]
- Calkins, M. W. Association. *Psychological Review*, 1894, 1, 476-483. [286]
- Calkins, M. W. Association. II. *Psychological Review*, 1896, 3, 32-49 [286]
- Camp, D. S., Raymond, G. A., and Church, R. M. Temporal relationship between response and punishment. *Journal of Experimental Psychology*, 1967, 74, 114-123. [100]
- Campbell, B. A., and Teghtsoonian, R. Electrical and behavioral effects of different types of shock stimuli on the rat. *Journal of Comparative and Physiological Psychology*, 1958, 51, 185-192. [100]
- Capaldi, E. J., and Stevenson, H. W. Response reversal following different amounts of training. *Journal of Comparative and Physiological Psychology*, 1957, 50, 195-198. [303]
- Carmichael, L. C., Hogan, H. P., and Walter, A. A. An experimental study of the effect of language on the reproduction of visually perceived form. *Journal of Experimental Psychology*, 1932, 15, 73-86. [330]
- Catania, A. C. Concurrent performances: a baseline for the study of reinforcement magnitude. *Journal of the Experimental Analysis of Behavior*. 1963, 6, 299-300. (a) [196]
- Catania, A. C. Concurrent performances: reinforcement interaction and response independence. *Journal of the Experimental Analysis of Behavior*. 1963, 6, 253-263. (b) [194, 195]
- Catania, A. C. Concurrent operants. In W. K. Honig (Ed.) *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. [192]
- Catania, A. C. Glossary. In A. C. Catania (Ed.) *Contemporary research in operant behavior*. Glenview, Ill.: Scott, Foresman, 1968. [75, 183]
- Catania, A. C. Concurrent performances: inhibition of one response by reinforcement of another. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 731-744. (a) [189, 194-195]
- Catania, A. C. On the vocabulary and the grammar of behavior. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 845-846. (b) [73]
- Catania, A. C. Reinforcement schedules and psychophysical judgments: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.) *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970. [127, 129]
- Catania, A. C. Reinforcement schedules: the role of responses preceding the one that produces the reinforcer. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 271-287. [176]

- Catania, A. C. Chomsky's formal analysis of natural languages: a behavioral translation. *Behaviorism*, 1972, 1, 1-15. [256, 356]
- Catania, A. C. The concept of the operant in the analysis of behavior. *Behaviorism*, 1973, 1, 103-116. (a) [127]
- Catania, A. C. The nature of learning. In J. A. Nevin and G. Reynolds (Eds.) *The study of behavior: learning, motivation, emotion, and instinct*. Glenview, Ill.: Scott, Foresman, 1973 (b) [42, 73]
- Catania, A. C. The psychologies of structure, function, and development. *American Psychologist*, 1973, 28, 434-443. (c) [11, 231, 356]
- Catania, A. C. Self-inhibiting effects of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 517-526. (d) [194]
- Catania, A. C. Freedom and knowledge: an experimental analysis of preference in pigeons. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 89-106. [13, 190, 196]
- Catania, A. C. The myth of self-reinforcement. *Behaviorism*, 1975, 3, 192-199. [350]
- Catania, A. C. Drug effects and concurrent performances. *Pharmacological Reviews*, 1976, 27, 385-394. [193]
- Catania, A. C. The psychology of learning: some lessons from the Darwinian revolution. *Annals of the New York Academy of Sciences*, 1978, 309, 18-28. [12, 357, 358]
- Catania, A. C., and Cutts, D. Experimental control of superstitious responding in humans. *Journal of the Experimental Analysis of Behavior*. 1963, 6, 203-208. [193]
- Catania, A. C., and Gill, C. A. Inhibition and behavioral contrast. *Psychonomic Science*, 1964, 1, 257-258. [188]
- Catania, A. C., Matthews, T. J., Silverman, P. J., and Yohalem, R. Yoked variable-ratio and variable-interval responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 1977, 28, 155-161. [174-175]
- Catania, A. C., and Reynolds, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 327-383. [169, 171, 180]
- Catania, A. C., Silverman, P. J., and Stubbs, D. A. Concurrent performances: stimulus-control gradients during schedules of signalled and unsignalled concurrent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 99-107. [153]
- Chang, J.-J., and Shepard, R. N. Meaningfulness in classification learning with pronounceable trigrams. *Journal of Verbal Learning and Verbal Behavior*, 1964, 3, 85-90. [294]
- Chomsky, N. Three models for the description of language. *IRE Transactions on Information Theory*, 1956, IT-2, 113-124. [256]
- Chomsky, N., and Miller, G. A. Introduction to the formal analysis of natural languages. In R. D. Luce, R. R. Bush, and E. Galanter (Eds.) *Handbook of mathematical psychology. Volume II*. New York: Wiley, 1963. [256, 259]
- Christen, F., and Bjork, R. A. On updating the loci in the Method of Loci. Paper presented at the Psychonomic Society, St. Louis, Mo., November, 1976. [311]
- Chung, S.-H. Effects of delayed reinforcement in a concurrent situation. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 439-444. [196]

- Chung, S.-H., and Herrnstein, R. J. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 67-74. [196]
- Church, R. M. The varied effects of punishment on behavior. *Psychological Review*, 1963, 70, 369-402. [97]
- Church, R. M. Response suppression. In B. A. Campbell and R. M. Church (Eds.) *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969. [100]
- Cieutat, V. J., Stockwell, F. E., and Noble, C. E. The interaction of ability and amount of practice with stimulus and response meaningfulness (m , m') in paired-associate learning. *Journal of Experimental Psychology*, 1958, 56, 193-202. [286]
- Collins, A. M., and Quillian, M. R. Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 1969, 8, 240-247. [271-272]
- Collins, A. M., and Quillian, M. R. Does category size affect categorization time? *Journal of Verbal Learning and Verbal Behavior*, 1970, 9, 432-438. [272]
- The compact edition of the Oxford English dictionary*. Oxford: Oxford University Press, 1971. [etymologies]
- Conrad, C. Cognitive economy in semantic memory. *Journal of Experimental Psychology*, 1972, 92, 149-154. [272]
- Conrad, R. Acoustic confusions in immediate memory. *British Journal of Psychology*, 1964, 55, 75-84. [327]
- Craik, F. I. M., and Lockhart, R. S. Levels of processing: a framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 1972, 11, 671-684. [329]
- Crowder, R. G. *Principles of learning and memory*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1976. [318, 323, 344]
- Cumming, W. W., and Berryman, R. The complex discriminated operant: studies of matching-to-sample and related problems. In D. I. Mostofsky (Ed.) *Stimulus generalization*. Stanford, Calif.: Stanford University Press, 1965. [156]
- Dalrymple-Alford, E. C. Response bias and judgments of the location of clicks in sentences. *Perception and Psychophysics*, 1976, 19, 303-308. [262]
- D'Amato, M. R., and Fazzaro, J. Attention and cue-producing behavior in the monkey. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 469-473. [157]
- D'Amato, M. R., and Jagoda, H. Analysis of the role of overlearning in discrimination reversal. *Journal of Experimental Psychology*, 1961, 61, 45-50. [303]
- Daniels, A. H. The memory after-image and attention. *American Journal of Psychology*, 1895, 6, 558-564. [316]
- Danks, J. H., and Glucksberg, S. Psychological scaling of linguistic properties. *Language and Speech*, 1970, 13, 118-140. [261]
- Darwin, C. *On the origin of species*. London: John Murray, 1859 (reprinted Cambridge, Mass.: Harvard University Press, 1966). [358-360]
- Darwin, C. J., Turvey, M. T., and Crowder, R. G. An auditory analogue of the Sperling partial report procedure: evidence for brief auditory storage. *Cognitive Psychology*, 1972, 3, 255-267. [316]
- Davis, H., and McIntire, R. W. Conditioned suppression under positive, negative, and no contingency between conditioned and unconditioned stimuli. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 633-640. [219]

- Day, W. F. On certain similarities between the philosophical investigations of Ludwig Wittgenstein and the operationism of B. F. Skinner. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 489-506. [241]
- Deese, J. Influence of inter-item associative strength upon immediate free recall. *Psychological Reports*, 1959, 5, 305-312. [292]
- Dews, P. B. Free-operant behavior under conditions of delayed reinforcement: I. CRF-type schedules. *Journal of the Experimental Analysis of Behavior*, 1960, 3, 221-234. [179, 185]
- Dews, P. B. The effect of multiple S^{Δ} periods on responding on a fixed-interval schedule. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 369-374. [176, 181]
- Dews, P. B. Drugs in psychology. A commentary on Travis Thompson and Charles R. Schuster's *Behavioral Pharmacology*. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 395-406. (a) [218]
- Dews, P. B. The theory of fixed-interval responding. In W. N. Schoenfeld (Ed.) *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970. (b) [181]
- Dill, L. M. The escape response of the zebra danio (*Brachydanio rerio*) II. The effect of experience. *Animal Behaviour*, 1974, 22, 723-730. [55]
- Dinsmoor, J. A. Punishment: I. The avoidance hypothesis. *Psychological Review*, 1954, 61, 34-46. [114]
- Duncker, K. On problem solving. *Psychological Monographs*, 1945, 58 (5), Whole No. 270. [350]
- Dunham, P. J. The nature of reinforcing stimuli. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [78]
- Ebbinghaus, H. *Über das Gedächtnis*. Leipzig: Duncker and Humblot, 1885 (*Memory*, tr. by H. A. Ruger and C. E. Bussenius, New York: Teachers College, 1913; reprinted by Dover, 1964). [243, 282-285, 320-321]
- Eckerman, D. A., and Lanson, R. N. Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement, and extinction. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 73-80. [123]
- Edgerton, V. R. Exercise and the growth and development of muscle tissue. In G. L. Rarick (Ed.) *Physical activity. Human growth and development*. New York: Academic Press, 1973. [6]
- Egeth, H. Selective attention. *Psychological Bulletin*, 1967, 67, 41-57. [144]
- Egger, M. D., and Miller, N. E. Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology*, 1962, 64, 97-104. [190]
- Eimas, P. D., Siqueland, E. R., Jusczyk, P., and Vigorito, J. Speech perception in early infancy. *Science*, 1971, 171, 303-306. [232]
- Ekstrand, B. R., Wallace, W. P., and Underwood, B. J. A frequency theory of verbal-discrimination learning. *Psychological Review*, 1966, 73, 566-578. [293]
- Emmons, W. H., and Simon, C. W. The non-recall of material presented during sleep. *American Journal of Psychology*, 1956, 69, 76-81. [321]
- Esper, E. A. A contribution to the experimental study of analogy. *Psychological Review*, 1918, 25, 468-487. [241, 274]

- Estes, W. K. An experimental study of punishment. *Psychological Monographs*, 1944, 57, No. 263. [96]
- Estes, W. K. The problem of inference from curves based on group data. *Psychological Bulletin*, 1956, 53, 134-140. [67]
- Estes, W. K. All-or-none processes in learning and retention. *American Psychologist*, 1964, 19, 16-25. [287]
- Estes, W. K., and Skinner, B. F. Some quantitative properties of anxiety. *Journal of Experimental Psychology*, 1941, 29, 390-400. [215, 217]
- Falk, J. L. Production of polydipsia in normal rats by an intermittent food schedule. *Science*, 1961, 133, 195-196. [48]
- Falk, J. L. The nature and determinants of adjunctive behavior. *Physiology and Behavior*, 1971, 6, 577-588. [48]
- Fantino, E. Conditioned reinforcement: choice and information. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [190, 196]
- Farmer, J. Properties of behavior under random interval reinforcement schedules. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 607-616. [169]
- Fearing, F. *Reflex action. A study in the history of physiological psychology (1930)*. New York: Hafner, 1964. [40]
- Felton, M., and Lyon, D. O. The post-reinforcement pause. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 131-134. [182]
- Ferster, C. B. Sustained behavior under delayed reinforcement. *Journal of Experimental Psychology*, 1953, 45, 218-224. [179]
- Ferster, C. B. Control of behavior in chimpanzees and pigeons by time out from positive reinforcement. *Psychological Monographs*, 1958, 72 (8, Whole No. 461). [113]
- Ferster, C. B. Intermittent reinforcement of matching to sample in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1960, 3, 259-272. [156]
- Ferster, C. B., and Hammer, C. E., Jr. Synthesizing the components of arithmetic behavior. In W. K. Honig (Ed.) *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. [253]
- Ferster, C. B., and Skinner, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957. [75, 130, 168, 174, 196]
- Fillenbaum, S. Memory for gist: some relevant variables. *Language and Speech*, 1966, 9, 217-227. [269]
- Fillenbaum, S., and Rapoport, A. *Structures in the subjective lexicon*. New York: Academic Press, 1971. [271]
- Findley, J. D. An experimental outline for building and exploring multi-operant behavior repertoires. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 113-166. [183, 192]
- Findley, J. D., and Brady, J. V. Facilitation of large ratio performance by use of conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 125-129. [192]
- Fischer, K. W. Development of intelligence: Piaget's theory. In K. W. Fischer, A. Lazerson, and P. Shaver (Eds.) *Psychology today: an introduction* (3rd ed.). Del Mar, Calif.: CRM Books, 1975. [347]
- Fisher, R. P., and Craik, F. I. M. Interaction between encoding and retrieval operations in cued recall. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, 3, 701-711. [331, 332]

- Fitts, P. M. Perceptual-motor skill learning. In A. W. Melton (Ed.) *Categories of human learning*. New York: Academic Press, 1964. [132]
- Flavell, J. H., Friedrichs, A. G., and Hoyt, J. D. Developmental changes in memorization processes. *Cognitive Psychology*, 1970, 1, 324-340. [324]
- Fleshler, M., and Hoffman, H. S. A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 529-530. [169]
- Fodor, J. A., and Bever, T. G. The psychological reality of linguistic segments. *Journal of Verbal Learning and Verbal Behavior*, 1965, 4, 414-420. [261]
- Fowler, H., and Trapold, M. A. Escape performance as a function of delay of reinforcement. *Journal of Experimental Psychology*, 1962, 63, 464-467. [105-106]
- Galton, F. Psychometric experiments. *Brain*, 1879, 2, 149-162. [242, 270]
- Gamzu, E. R., and Williams, D. R. Associative factors underlying the pigeon's key pecking in auto-shaping procedures. *Journal of the Experimental Analysis of Behavior*, 1973, 19, 225-232. [214]
- Garcia, J., and Koelling, R. A. Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 1966, 4, 123-124. [160-161]
- Gardner, E. T., and Lewis, P. Negative reinforcement with shock-frequency increase. *Journal of the Experimental Analysis of Behavior*, 1976, 25, 3-14. [111]
- Gardner, R. A., and Gardner, B. T. Teaching sign language to a chimpanzee. *Science*, 1969, 165, 664-672. [248, 274]
- Garner, W. R. *The processing of information and structure*. Potomac, Md.: Lawrence Erlbaum Associates, 1974. [346]
- Garrett, M., Bever, T. G., and Fodor, J. The active use of grammar in speech perception. *Perception and Psychophysics*, 1966, 1, 30-32. [262]
- Geller, I. The acquisition and extinction of conditioned suppression as a function of the base-line reinforcer. *Journal of the Experimental Analysis of Behavior*, 1960, 3, 235-240. [215-216]
- Gibson, E. J. Learning to read. *Science*, 1965, 148, 1066-1072. [156, 233, 303, 327]
- Gilbert, R. M. Variation and selection of behavior. In R. M. Gilbert and J. R. Millenson (Eds.) *Reinforcement: behavioral analyses*. New York: Academic Press, 1972. [358]
- Ginsburg, H., and Oppen, S. *Piaget's theory of intellectual development. An introduction*. Englewood Cliffs, N.J.: Prentice-Hall, 1969. [346]
- Girgus, J. S. A developmental study of the effect of eye movements on the processing of shape information. *Journal of Experimental Child Psychology*, 1976, 22, 386-399. [91]
- Glanzer, M. Short-term storage and long-term storage in recall. *Journal of Psychiatric Research*, 1971, 8, 423-438. [290]
- Glanzer, M., and Cunitz, A. R. Two storage mechanisms in free recall. *Journal of Verbal Learning and Verbal Behavior*, 1966, 5, 351-360. [290]
- Glanzer, M., and Dolinsky, R. The anchor for the serial position curve. *Journal of Verbal Learning and Verbal Behavior*, 1965, 4, 267-273. [283]
- Glanzer, M., and Koppenaal, L. The effect of encoding tasks on free recall: stages and levels. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 21-28. [329]

- Glaze, J. A. The association value of non-sense syllables. *Journal of Genetic Psychology*, 1928, 35, 255-267. [270, 282]
- Gleitman, L. R., and Gleitman, H. *Phrase and paraphrase*. New York: Norton, 1970. [261, 263]
- Gleitman, L. R., and Rozin, P. Teaching reading by use of a syllabary. *Reading Research Quarterly*, 1973, 8, 447-483. [234, 295]
- Glickman, S. E. Perseverative neural processes and consolidation of the memory trace. *Psychological Bulletin*, 1961, 58, 218-233. [322]
- Glucksberg, S., and Danks, J. H. Grammatical structure and recall: A function of the space in immediate memory or of recall delay? *Perception and Psychophysics*, 1969, 6, 113-117. [265]
- Glucksberg, S., and Danks, J. H. *Experimental psycholinguistics*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1975. [262]
- Goldiamond, I. Perception, language, and conceptualization rules. In B. Kleinmuntz (Ed.) *Problem solving*. New York: Wiley, 1966. [294-295]
- Goldstein, H., Krantz, D. L., and Rains, J. D. *Controversial issues in learning*. New York: Appleton-Century-Crofts, 1965. [65]
- Gollub, L. R. The relations among measures of performance on fixed-interval schedules. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 337-343. [181]
- Gollub, L. R. Stimulus generalization of response-position in the rat. *Psychonomic Science*, 1966, 6, 433-434. [123]
- Gollub, L. R. Conditioned reinforcement: schedule effects. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [192]
- Gormezano, I. Investigations of defense and reward conditioning in the rabbit. In A. H. Black and W. F. Prokasy (Eds.) *Classical conditioning II: current research and theory*. New York: Appleton-Century-Crofts, 1972. [206]
- Gould, S. J. Darwin's "Big Book." *Science*, 1975, 188, 824-827. [359]
- Green, L., and Rachlin, H. Pigeon's preferences for stimulus information: effects of amount of information. *Journal of the Experimental Analysis of Behavior*, 1977, 27, 255-263. [196]
- Greeno, J. G. Hobbits and orcs: acquisition of a sequential concept. *Cognitive Psychology*, 1974, 6, 270-292. [348]
- Greenspoon, J. The reinforcing effect of two spoken sounds on the frequency of two responses. *American Journal of Psychology*, 1955, 68, 409-416. [244]
- Groninger, L. D. Predicting recognition during storage: the capacity of the memory system to evaluate itself. *Bulletin of the Psychonomic Society*, 1976, 7, 425-428. [325]
- Groves, P. M., and Thompson, R. F. Habituation: a dual-process theory. *Psychological Review*, 1970, 77, 419-450. [52]
- Guess, D. A functional analysis of receptive language and productive speech: acquisition of the plural morpheme. *Journal of Applied Behavior Analysis*, 1969, 2, 55-64. [275]
- Guthrie, E. R. *The psychology of learning*. New York: Harper, 1935. [84, 204]
- Guthrie, E. R., and Horton, G. P. *Cats in a puzzle box*. New York: Rinehart, 1946. [84]
- Guttman, N. The pigeon and the spectrum and other perplexities. *Psychological Reports*, 1956, 2, 449-460. [151]

- Guttman, N. Generalization gradients around stimuli associated with different reinforcement schedules. *Journal of Experimental Psychology*, 1959, 58, 335-340. [149]
- Guttman, N., and Kalish, H. I. Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 1956, 51, 79-88. [150]
- Haber, R. N. Eidetic images. *Scientific American*, 1969, 220, 36-44. [341]
- Hadamard, J. *The psychology of invention in the mathematical field*. Princeton: Princeton University Press, 1949. [330]
- Hailman, J. P. How an instinct is learned. *Scientific American*, 1969, 221 (6), 98-106. [29]
- Hanson, H. M. Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology*, 1959, 58, 321-334. [149-150]
- Harlow, H. F. The formation of learning sets. *Psychological Review*, 1949, 56, 51-65. [162-163]
- Harris, C. S. Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological Review*, 1965, 72, 419-444. [90]
- Hart, J. T. Second-try recall, recognition, and the memory-monitoring process. *Journal of Educational Psychology*, 1967, 58, 193-197. [325]
- Harzem, P., Lee, I., and Miles, T. R. The effects of pictures on learning to read. *British Journal of Educational Psychology*, 1976, 46, 318-322. [356]
- Hawkes, L., and Shimp, C. P. Reinforcement of behavioral patterns: shaping a scallop. *Journal of the Experimental Analysis of Behavior*, 1975, 23, 3-16. [181]
- Haygood, R. C., and Bourne, L. E. Attribute and rule learning aspects of conceptual behavior. *Psychological Review*, 1965, 72, 175-195. [346]
- Healy, A. F. Detection errors on the word *the*: evidence for reading units larger than letters. *Journal of Experimental Psychology: Human Perception and Performance*, 1976, 2, 235-242. [293]
- Hearst, E. The behavioral effects of some temporally defined schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 45-55. [177]
- Hearst, E., Besley, S., and Farthing, G. W. Inhibition and the stimulus control of operant behavior. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 373-409. [149, 151]
- Hearst, E., and Jenkins, H. M. *Sign-tracking: the stimulus-reinforcer relation and directed action*. Austin, Texas: Psychonomic Society, 1974. [213, 214]
- Hearst, E., Koresko, M. B., and Poppen, R. Stimulus generalization and the response-reinforcement contingency. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 369-379. [147]
- Hefferline, R. F., and Keenan, B. Amplitude-induction gradient of a small-scale (covert) operant. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 307-315. [127-128]
- Hein, H. The endurance of the mechanism-vitalism controversy. *Journal of the History of Biology*, 1972, 5, 159-188. [357]
- Heinemann, E. G., and Avin, E. On the development of stimulus control. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 183-195. [159]
- Held, R., and Hein, A. Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 1963, 56, 872-876. [89-90]

- Hemmes, N. S. Behavioral contrast in pigeons depends upon the operant. *Journal of Comparative and Physiological Psychology*, 1973, 85, 171-178. [188]
- Hendry, D. P. Reinforcing value of information: fixed-ratio schedules. In D. P. Hendry (Ed.) *Conditioned reinforcement*. Homewood, Ill.: Dorsey, 1969. [190]
- Hendry, D. P., and Hendry, L. S. Partial negative reinforcement: fixed-ratio escape. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 519-523. [185]
- Herrick, R. M. The successive differentiation of a lever displacement response. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 211-215. [125]
- Herrick, R. M., Myers, J. L., and Korotkin, A. L. Changes in S^D and in S^A rates during the development of an operant discrimination. *Journal of Comparative and Physiological Psychology*, 1959, 52, 359-364. [140]
- Herrnstein, R. J. Some factors influencing behavior in a two-response situation. *Transactions of the New York Academy of Sciences*, 1958, 21, 35-45. [192]
- Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 267-272. (a) [193]
- Herrnstein, R. J. Stereotypy and intermittent reinforcement. *Science*, 1961, 133, 2067-2069. (b) [121]
- Herrnstein, R. J. Aperiodicity as a factor in choice. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 179-182. (a) [196]
- Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 27-36. (b). [196]
- Herrnstein, R. J. Superstition: a corollary of the principles of operant conditioning. In W. K. Honig (Ed.) *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. [84, 177-178]
- Herrnstein, R. J. Method and theory in the study of avoidance. *Psychological Review*, 1969, 76, 49-69. [111]
- Herrnstein, R. J. On the law of effect. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 243-266. [188, 193]
- Herrnstein, R. J., and Himeline, P. N. Negative reinforcement as shock-frequency reduction. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 421-430. [111]
- Herrnstein, R. J., and Loveland, D. H. Complex visual concept in the pigeon. *Science*, 1964, 146, 549-551. [159]
- Herrnstein, R. J., and Loveland, D. H. Maximizing and matching on concurrent ratio schedules. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 107-116. [193]
- Herrnstein, R. J., Loveland, D. H., and Cable, C. Natural concepts in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 1976, 2, 285-311. [159-160]
- Herrnstein, R. J., and Morse, W. H. Some effects of response-independent positive reinforcement on maintained operant behavior. *Journal of Comparative and Physiological Psychology*, 1957, 50, 461-467. [217]
- Herrnstein, R. J., and Morse, W. H. A conjunctive schedule of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 15-24. [185-186]
- Herrnstein, R. J., and Sidman, M. Avoidance conditioning as a factor in the effects of unavoidable shocks on food-reinforced behavior. *Journal of Comparative and Physiological Psychology*, 1958, 51, 380-385. [217]

References

- Hess, E. H. *Imprinting*. New York: van Nostrand Reinhold, 1973. [59]
- Higbee, K. L. *Your memory: how it works and how to improve it*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [313]
- Hilgard, E. R. Method and procedures in the study of learning. In S. S. Stevens (Ed.) *Handbook of experimental psychology*. New York: Wiley, 1951. [65]
- Hineline, P. N. Negative reinforcement without shock reduction. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 259-268. [111]
- Hineline, P. N. Negative reinforcement and avoidance. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [111, 113]
- Hively, W. Programming stimuli in matching to sample. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 279-298. [154]
- Hoffman, H. S., and Fleshler, M. Aversive control with the pigeon. *Journal of the Experimental Analysis of Behavior*, 1959, 2, 213-218. [106]
- Hoffman, H. S., and Fleshler, M. The course of emotionality in the development of avoidance. *Journal of Experimental Psychology*, 1962, 64, 288-294. [111, 221]
- Hoffman, H. S., and Ratner, A. M. A reinforcement model of imprinting: implications for socialization in monkeys and men. *Psychological Review*, 1973, 80, 527-544. [59]
- Hogan, J. A. The development of a hunger system in young chicks. *Behaviour*, 1971, 39, 128-201. [56]
- Holland, J. G. Human vigilance. *Science*, 1958, 128, 61-67. [180]
- Holmes, V. M., and Forster, K. I. Click location and syntactic structure. *Perception and Psychophysics*, 1972, 12, 9-15. [262]
- Holz, W. C., and Azrin, N. H. Discriminative properties of punishment. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 225-232. [102-103]
- Honig, W. K. Generalization of extinction on the spectral continuum. *Psychological Record*, 1961, 11, 269-278. [151]
- Honig, W. K. Prediction of preference, transposition, and transposition-reversal from the generalization gradient. *Journal of Experimental Psychology*, 1962, 64, 239-248. [304]
- Honig, W. K., Boneau, C. A., Burstein, K. R., and Pennypacker, H. S. Positive and negative generalization gradients obtained after equivalent training conditions. *Journal of Comparative and Physiological Psychology*, 1963, 56, 111-116. [151-152]
- Hull, C. L. Quantitative aspects of the evolution of concepts: an experimental study. *Psychological Monographs*, 1920, 28, No. 123. [238]
- Hull, C. L. The factor of the conditioned reflex. In C. Murchison (Ed.) *Handbook of general experimental psychology*. Worcester, Mass.: Clark University Press, 1934. [205]
- Hull, C. L. *Principles of behavior*. New York: Appleton-Century-Crofts, 1943. [18, 108, 204]
- Hunter, W. S. The behavior of raccoons in a double-alternation temporal maze. *Journal of Genetic Psychology*, 1928, 35, 374-388. [132]
- Hutchinson, R. R., Renfrew, J. W., and Young, G. A. Effects of long-term shock and associated stimuli on aggressive and manual responses. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 141-166. [49, 53, 57]

- Huttenlocher, J. Constructing spatial images: a strategy in reasoning. *Psychological Review*, 1968, 75, 550-560. [341]
- Huttenlocher, J., and Strauss, S. Comprehension and a statement's relation to the situation it describes. *Journal of Verbal Learning and Verbal Behavior*, 1968, 7, 300-304. [341]
- Jacobs, J. Experiments on "prehension." *Mind*, 1887, 12, 75-79. [316]
- Jans, J. E. *Short-term remembering of discriminative stimuli in pigeons*. Unpublished Masters thesis, University of Maryland Baltimore County, 1976. [310]
- Jarvella, R. J. Syntactic processing of connected speech. *Journal of Verbal Learning and Verbal Behavior*, 1971, 10, 409-416. [261]
- Jaynes, J. *The origin of consciousness in the breakdown of the bicameral mind*. Boston: Houghton Mifflin, 1977. [241, 275]
- Jenkins, H. M. Generalization gradients and the concept of inhibition. In D. I. Mostofsky (Ed.) *Stimulus generalization*. Stanford, Calif.: Stanford University Press, 1965. [151]
- Jenkins, H. M. Sequential organization in schedules of reinforcement. In W. N. Schoenfeld (Ed.) *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970. [176]
- Jenkins, H. M., and Harrison, R. H. Effect of discrimination training on auditory generalization. *Journal of Experimental Psychology*, 1960, 59, 246-253. [147-148]
- Jenkins, H. M., and Harrison, R. H. Generalization gradients of inhibition following auditory discrimination learning. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 435-441. [151]
- Jenkins, H. M., and Moore, B. R. The form of the auto-shaped response with food or water reinforcers. *Journal of the Experimental Analysis of Behavior*, 1973, 20, 163-181. [214]
- Jenkins, H. M., and Sainsbury, R. S. Discrimination learning with the distinctive feature on positive or negative trials. In D. I. Mostofsky (Ed.) *Attention: contemporary theory and analysis*. New York: Appleton-Century-Crofts, 1970. [153]
- Jenkins, J. G. Instruction as a factor in "incidental" learning. *American Journal of Psychology*, 1933, 45, 471-477. [278]
- Jenkins, J. G., and Dallenbach, K. M. Obliviscence during sleep and waking. *American Journal of Psychology*, 1924, 35, 605-612. [322]
- Jenkins, W. O., Pascal, G. R., and Walker, R. W. Deprivation and generalization. *Journal of Experimental Psychology*, 1958, 56, 274-277. [147]
- Jensen, D. D. Experiments on learning in paramecia. *Science*, 1957, 125, 191-192. [5]
- Johnson, N. F. The psychological reality of phrase-structure rules. *Journal of Verbal Learning and Verbal Behavior*, 1965, 4, 469-475. [261-262]
- Jonçich, G. *The sane positivist. A biography of Edward L. Thorndike*. Middletown, Conn.: Wesleyan University Press, 1968. [298]
- Jost, A. Die Associationsfestigkeit in ihrer Abhängigkeit von der Verteilung der Wiederholungen. *Zeitschrift für Psychologie*, 1897, 14, 436-472. [321]
- Kamin, L. J. The effects of termination of the CS and avoidance of the US on avoidance learning. *Journal of Comparative and Physiological Psychology*, 1956, 49, 420-424. [111]

- Kamin, L. J. The effects of termination of the CS and avoidance of the US on avoidance learning: an extension. *Canadian Journal of Psychology*, 1957, 11, 48-56. [111]
- Kamin, L. J. The retention of an incompletely learned avoidance response. *Journal of Comparative and Physiological Psychology*, 1957, 50, 457-460. [322]
- Kamin, L. J. Predictability, surprise, attention and conditioning. In B. A. Campbell and R. M. Church (Eds.) *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969. [210]
- Katz, S., and Gruenwald, P. The abstraction of linguistic ideas in "meaningless" sentences. *Memory and Cognition*, 1974, 2, 737-741. [269]
- Kaufman, A., and Baron, A. Suppression of behavior by timeout punishment when suppression results in loss of positive reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 595-607. [113]
- Kausler, D. H., and Boka, J. A. Effects of double functioning on verbal-discrimination learning. *Journal of Experimental Psychology*, 1968, 76, 558-567. [294]
- Kelleher, R. T. Concept formation in chimpanzees. *Science*, 1958, 128, 777-778. (a) [159]
- Kelleher, R. T. Stimulus-producing responses in chimpanzees. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 87-102. (b) [157]
- Kelleher, R. T. Chaining and conditioned reinforcement. In W. K. Honig (Ed.) *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. [192]
- Kelleher, R. T., and Fry, W. T. Stimulus functions in chained fixed-interval schedules. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 167-173. [192]
- Kelleher, R. T., and Gollub, L. R. A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 543-597. [8, 183, 190]
- Kelleher, R. T., and Morse, W. H. Schedules using noxious stimuli. III. Responding maintained with response-produced electric shocks. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 819-838. [101]
- Kelleher, R. T., Riddle, W. C., and Cook, L. Observing responses in pigeons. *Journal of the Experimental Analysis of Behavior*, 1962, 5, 3-13. [157]
- Kelleher, R. T., Riddle, W. C., and Cook, L. Persistent behavior maintained by unavoidable shocks. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 507-517. [217]
- Keller, F. S. Light aversion in the white rat. *Psychological Record*, 1941, 4, 235-250. [105, 107]
- Keller, F. S. The phantom plateau. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 1-13. [327]
- Keller, F. S., and Schoenfeld, W. N. *Principles of psychology*. New York: Appleton-Century-Crofts, 1950. [159]
- Keller, K. The role of elicited responding in behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 249-257. [190, 214]
- Kendall, S. B. Spontaneous recovery after extinction with periodic time-outs. *Psychonomic Science*, 1965, 2, 117-118. [81]
- Kendler, H. H., and Kendler, T. S. Vertical and horizontal processes in problem solving. *Psychological Review*, 1962, 69, 1-16. [304]

- Keppel, G., and Underwood, B. J. Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior*, 1962, 1, 153-156. [319]
- Killeen, P. A yoked-chamber comparison of concurrent and multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1972, 18, 13-22. [196]
- Killeen, P. On the temporal control of behavior. *Psychological Review*, 1975, 82, 89-115. [84, 181]
- Kimble, G. A. Conditioning as a function of the time between conditioned and unconditioned stimuli. *Journal of Experimental Psychology*, 1947, 37, 1-15. [202]
- Kimble, G. A. *Hilgard and Marquis' Conditioning and learning*. (Second edition) New York: Appleton-Century-Crofts, 1961. [4, 81]
- Kimble, D. P., and Ray, R. S. Reflex habituation and potentiation in *Rana pipiens*. *Animal Behaviour*, 1965, 13, 530-533. [53]
- Kimmel, H. D. Notes from 'Pavlov's Wednesdays': Pavlov's law of effect. *American Journal of Psychology*, 1976, 89, 553-556. [207]
- Kincaid, J. P., and Wickens, D. D. Temporal gradient of release from proactive inhibition. *Journal of Experimental Psychology*, 1970, 86, 313-316. [319]
- Kintsch, W. An experimental analysis of single stimulus tests and multiple-choice tests of recognition memory. *Journal of Experimental Psychology*, 1968, 76, 1-6. [296]
- Kintsch, W. *The representation of meaning in memory*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1974. [269, 272]
- Kish, G. B. Studies of sensory reinforcement. In W. K. Honig (Ed.) *Operant behavior: Areas of research and application*. New York: Appleton-Century-Crofts, 1966. [89]
- Köhler, W. *The mentality of apes*. (tr. by E. Winter; Second revised edition) London: Routledge & Kegan Paul, 1927. [25-26, 29, 351]
- Köhler, W. *Gestalt psychology*. New York: Liveright, 1929. [330]
- Kolers, P. A. Reading and talking bilingually. *American Journal of Psychology*, 1966, 79, 357-376. [272, 288, 335]
- Kolers, P. A., and Paley, R. Knowing not. *Memory and Cognition*, 1976, 4, 553-558. [325, 333]
- Konorski, J. *Conditioned reflexes and neuron organization*. New York: Cambridge University Press, 1948. [204]
- Krechevsky, I. 'Hypotheses' in rats. *Psychological Review*, 1932, 39, 516-532. [133]
- Kuch, D. O., and Platt, J. R. Reinforcement rate and interresponse time differentiation. *Journal of the Experimental Analysis of Behavior*, 1976, 26, 471-486. [175]
- Landauer, T. K., and Meyer, D. E. Category size and semantic-memory retrieval. *Journal of Verbal Learning and Verbal Behavior*, 1972, 11, 539-549. [272]
- Lander, D. G., and Irwin, R. J. Multiple schedules: effects of the distribution of reinforcements between components on the distribution of responses between components. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 517-524. [188]
- Lane, H. The motor theory of speech perception: a critical review. *Psychological Review*, 1965, 72, 275-309. [232]

- Lashley, K. S. The mechanism of vision. I. A method for rapid analysis of pattern vision in the rat. *Journal of Genetic Psychology*, 1930, 37, 453-460. [33-34]
- Lashley, K. S. The problem of serial order in behavior. In L. A. Jeffress (Ed.) *Cerebral mechanisms in behavior*. New York: Wiley, 1951. [134-135]
- Lattal, K. A. Combinations of response-reinforcer dependence and independence. *Journal of the Experimental Analysis of Behavior*, 1974, 22, 357-362. [178-179]
- Lawrence, D. H. Acquired distinctiveness of cues: transfer between discriminations on the basis of familiarity with the stimulus. I. *Journal of Experimental Psychology*, 1949, 39, 770-784. [145]
- Lawrence, D. H. Acquired distinctiveness of cues: transfer between discriminations on the basis of familiarity with the stimulus. II. *Journal of Experimental Psychology*, 1950, 40, 175-188. [145]
- Leitenberg, H. Conditioned acceleration and conditioned suppression in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 205-212. [217]
- Lenneberg, E. H. Understanding language without ability to speak: a case report. *Journal of Abnormal and Social Psychology*, 1962, 65, 419-425. [268]
- Lenneberg, E. H. *Biological foundations of language*. New York: Wiley, 1967. [273]
- Levine, M. Hypothesis behavior by humans during discrimination learning. *Journal of Experimental Psychology*, 1966, 71, 331-338. [133]
- Levitsky, D., and Collier, G. Schedule-induced wheel running. *Physiology and Behavior*, 1968, 3, 571-573. [49]
- Lieberman, A. M. The grammars of speech and language. *Cognitive psychology*, 1970, 1, 301-323. [232]
- Linden, E. *Apes, men and language*. New York: Penguin, 1976. [274]
- Loisette, A. *Assimilative memory, or, How to attend and never forget*. New York: Funk and Wagnalls, 1899. [312]
- Lorenz, K. The companion in the bird's world. *Auk*, 1937, 54, 245-273. [59]
- Luchins, A. S. Mechanization in problem solving. *Psychological Monographs*, 1942, 54 (6), Whole No. 248. [348-349]
- Luchins, A. S., and Luchins, E. H. New experimental attempts at preventing mechanization in problem solving. *Journal of General Psychology*, 1950, 42, 279-297. [348]
- MacCorquodale, K. On Chomsky's review of Skinner's *Verbal behavior*. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 83-99. [229]
- Mackintosh, N. J. Selective attention in animal discrimination learning. *Psychological Bulletin*, 1965, 64, 124-150. [144]
- Mackintosh, N. J. Stimulus control: attention factors. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N.J.: Prentice-Hall, 1977. [159]
- MacLeod, C. M. Bilingual episodic memory: acquisition and forgetting. *Journal of Verbal Learning and Verbal Behavior*, 1976, 15, 347-364. [335]
- Macphail, E. M. Avoidance responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 629-632. [110]
- Maier, N. R. F., Glaser, N. M., and Klee, J. B. Studies of abnormal behavior in the rat. III. The development of behavior fixations through frustration. *Journal of Experimental Psychology*, 1940, 26, 521-546. [303]

- Maier, S. F., Seligman, M. E. P., and Solomon, R. L. Pavlovian fear conditioning and learned helplessness: effects on escape and avoidance behavior of (a) the CS-US contingency and (b) the independence of the US and voluntary responding. In B. A. Campbell and R. M. Church (Eds.) *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969. [164]
- Malcolm, N. The myth of cognitive processes and structures. In T. Mischel (Ed.) *Cognitive development and epistemology*. New York: Academic Press, 1971. [338]
- Malott, R. W., and Cumming, W. W. Schedules of interresponse time reinforcement. *Psychological Record*, 1964, 14, 221-252. [130]
- Martin, E. Relation between stimulus recognition and paired-associate learning. *Journal of Experimental Psychology*, 1967, 74, 500-505. [296]
- Matthews, B. A., Shimoff, E., Catania, A. C., and Sagvolden, T. Uninstructed human responding: sensitivity to ratio and interval contingencies. *Journal of the Experimental Analysis of Behavior*, 1977, 27, 453-467. [180, 247]
- Matthews, W. A. Transformational complexity and short-term recall. *Language and Speech*, 1968, 11, 120-128. [265]
- Max, L. W. An experimental study of the motor theory of consciousness. I. History and critique. *Journal of General Psychology*, 1934, 11, 112-125. [235]
- McCoon, G. Organization of information in text memory. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 247-260. [335]
- McGaugh, J. L., and Dawson, R. G. Modification of memory storage processes. *Behavioral Science*, 1971, 16, 45-63. [322]
- McGeoch, J. A. *The psychology of human learning*. New York: Longmans, Green, 1942. [277, 282, 283, 298, 322]
- McNeill, D. *The acquisition of language*. New York: Harper and Row, 1970. [274]
- McReynolds, P. Statues, clocks, and computers: on the history of models in psychology. *Proceedings, 79th Annual Convention, American Psychological Association*, 1971, 715-716. [333]
- Mechner, F. A notation system for the description of behavioral procedures. *Journal of the Experimental Analysis of Behavior*, 1959, 2, 133-150. [185]
- Meehl, P. E. On the circularity of the Law of Effect. *Psychological Bulletin*, 1950, 47, 52-75. [75]
- Mehler, J. Some effects of grammatical transformations on the recall of English sentences. *Journal of Verbal Learning and Verbal Behavior*, 1963, 2, 346-351. [265, 266]
- Meisch, R. A., and Thompson, T. Ethanol intake in the absence of concurrent food reinforcement. *Psychopharmacologia*, 1971, 22, 72-79. [49]
- Melton, A. W. Implications of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 1963, 2, 1-21. [318-319, 320]
- Melton, A. W., and Martin, E. (Eds.) *Coding processes in human memory*. Washington, D. C.: Winston, 1972. [327]
- Meltzer, D., and Brahlek, J. A. Conditioned suppression and conditioned enhancement with the same positive UCS: an effect of CS duration. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 67-73. [219]
- Meltzoff, A. N., and Moore, M. K. Imitation of facial and manual gestures by human neonates. *Science*, 1977, 198, 75-78. [164]

- Michael, J. Positive and negative reinforcement, a distinction that is no longer necessary; or a better way to talk about bad things. *Behaviorism*, 1975, 3, 33-44. [104]
- Milgram, S. Behavioral study of obedience. *Journal of Abnormal and Social Psychology*, 1963, 67, 371-378. [247]
- Millenson, J. R. Random interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 437-443. [169]
- Miller, G. A. The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 1956, 63, 81-97. [314, 316, 327]
- Miller, G. A. Free recall of redundant strings of letters. *Journal of Experimental Psychology*, 1958, 56, 485-491. [283]
- Miller, G. A. Some psychological studies of grammar. *American Psychologist*, 1962, 17, 748-762. [265, 266, 271]
- Miller, G. A. Lexical memory. *Proceedings of the American Philosophical Society*, 1972, 116, 140-144. [334]
- Miller, G. A., and Isard, S. Some perceptual consequences of linguistic rules. *Journal of Verbal Learning and Verbal Behavior*, 1963, 2, 217-228. [266]
- Miller, G. A., and Selfridge, J. A. Verbal context and the recall of meaningful material. *American Journal of Psychology*, 1953, 63, 176-185. [257]
- Miller, N. E. Biofeedback and visceral learning. *Annual Review of Psychology*, 1978, 29, 373-404. [121]
- Miller, N. E., and Banuazizi, A. Instrumental learning by curarized rats of a specific-visceral response, intestinal or cardiac. *Journal of Comparative and Physiological Psychology*, 1968, 65, 1-7. [121-122]
- Miller, N. E., and Carmona, A. Modification of a visceral response, salivation in thirsty dogs, by instrumental training with water reward. *Journal of Comparative and Physiological Psychology*, 1967, 63, 1-6. [121, 208]
- Moeser, S. D. Inferential reasoning in episodic memory. *Journal of Verbal Learning and Verbal Behavior*, 1976, 15, 193-212. [335]
- Moeser, S. D., and Tarrant, B. L. Learning a network of comparisons. *Journal of Experimental Psychology: Human Learning and Memory*, 1977, 3, 643-659. [335, 341]
- Morgan, M. J., Fitch, M. D., Holman, J. G., and Lea, S. E. G. Pigeons learn the concept of an 'A'. *Perception*, 1976, 5, 57-66. [159]
- Morse, W. H. Intermittent reinforcement. In W. K. Honig (Ed.) *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. [79]
- Morse, W. H., and Kelleher, R. T. Determinants of reinforcement and punishment. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [101, 102, 192]
- Morse, W. H., Mead, R. N., and Kelleher, R. T. Modulation of elicited behavior by a fixed-interval schedule of electric shock presentation. *Science*, 1967, 157, 215-217. [51]
- Mowrer, O. H. On the psychology of "talking birds"—a contribution to language and personality theory. In O. H. Mowrer, *Learning theory and personality dynamics: selected papers*. New York: Ronald Press, 1950. [232]
- Mowrer, O. H. *Learning theory and behavior*. New York: Wiley, 1960. [204, 220]

- Mowrer, O. H., and Jones, H. M. Extinction and behavior variability as functions of effortfulness of task. *Journal of Experimental Psychology*, 1943, 33, 369-385. [83]
- Mowrer, O. H., and Lamoreaux, R. R. Fear as an intervening variable in avoidance conditioning. *Journal of Comparative Psychology*, 1946, 39, 29-50. [111]
- Moyer, R. S. Comparing objects in memory: evidence suggesting an internal psychophysics. *Perception and Psychophysics*, 1973, 13, 180-184. [341]
- Müller, G. E., and Pilzecker, A. Experimentelle Beiträge zur Lehre vom Gedächtnis. *Zeitschrift für Psychologie*, 1900, Ergänzungsband 1. [302]
- Murdock, B. B., Jr. The serial position effect in free recall. *Journal of Experimental Psychology*, 1962, 64, 482-488. [289]
- Murray, D. J. Research on human memory in the nineteenth century. *Canadian Journal of Psychology*, 1976, 30, 201-220. [309, 316]
- Myers, W. A. Observational learning in monkeys. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 225-235. [164]
- Neisser, U. *Cognitive psychology*. New York: Appleton-Century-Crofts, 1967. [293]
- Neuringer, A., and Schneider, B. A. Separating the effects of interreinforcement time and number of interreinforcement responses. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 661-667. [196]
- Nevin, J. A. Response strength in multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 389-408. [79, 131, 197]
- Nevin, J. A. and Shettleworth, S. J. An analysis of contrast effects in multiple schedules. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 305-315. [188]
- Newman, S. E. In search of associative symmetry. In C. P. Duncan, L. Sechrest, and A. W. Melton (Eds.) *Human memory: Festschrift in honor of Benton J. Underwood*. New York: Appleton-Century-Crofts, 1972. [288]
- Norman, D. A. Toward a theory of memory and attention. *Psychological Review*, 1968, 75, 522-536. [313-314]
- Norman, D. A. (Ed.) *Models of memory*. New York: Academic Press, 1970. [313-314, 325]
- Oden, G. C. Fuzziness in semantic memory: choosing exemplars of subjective categories. *Memory and Cognition*, 1977, 5, 198-204. [272]
- Olds, J., and Milner, P. Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of Comparative and Physiological Psychology*, 1954, 47, 419-427. [121]
- Orne, M. T. On the social psychology of the psychological experiment: with particular reference to demand characteristics and their implications. *American Psychologist*, 1962, 17, 776-783. [278]
- Osgood, C. E. The similarity paradox in human learning: a resolution. *Psychological Review*, 1949, 56, 132-143. [300-301]
- Osgood, C. E., Suci, G. J., and Tannenbaum, P. H. *The measurement of meaning*. Urbana: University of Illinois, 1957. [270]
- Overton, D. A. State-dependent or "dissociated" learning produced with phenobarbital. *Journal of Comparative and Physiological Psychology*, 1964, 57, 3-12. [333]
- The Oxford universal dictionary on historical principles* (Third edition, rev. with Addenda). Oxford: Clarendon, 1955. [etymologies]

- Paivio, A. Mental imagery in associative learning and memory. *Psychological Review*, 1969, 76, 241-263. [328]
- Paivio, A. *Imagery and verbal processes*. New York: Holt, Rinehart and Winston, 1971. [328, 340]
- Paivio, A. Neomentatism. *Canadian Journal of Psychology*, 1975, 29, 263-291. [338, 341]
- Paivio, A., and Madigan, S. A. Imagery and association value in paired associate learning. *Journal of Experimental Psychology*, 1968, 76, 35-39. [328]
- Palermo, D. S., and Eberhart, V. L. On the learning of morphological rules: an experimental analogy. *Journal of Verbal Learning and Verbal Behavior*, 1968, 7, 337-344. [275]
- Palermo, D. S., and Howe, H. E., Jr. An experimental analogy to the learning of past tense inflection rules. *Journal of Verbal Learning and Verbal Behavior*, 1970, 9, 410-416. [275]
- Partridge, E. *Origins. A short etymological dictionary of modern English*. (Second edition) New York: Macmillan, 1959. [etymologies]
- Pavlov, I. P. *Conditioned reflexes*. (tr. G. V. Anrep) London: Oxford University Press, 1927. [31-32, 40, 200]
- Peterson, L. R., and Peterson, M. J. Short-term retention of individual verbal items. *Journal of Experimental Psychology*, 1959, 58, 193-198. [316-317, 319]
- Peterson, N. Control of behavior by presentation of an imprinted stimulus. *Science*, 1960, 132, 1395-1396. [60]
- Peterson, N. Effect of monochromatic rearing on the control of responding by wavelength. *Science*, 1962, 136, 774-775. [151]
- Phillips, D. C. James, Dewey, and the reflex arc. *Journal of the History of Ideas*, 1971, 32, 555-568. [40]
- Piaget, J., and Inhelder, B. *The psychology of the child* (tr. by H. Weaver). New York: Basic Books, 1969. [346]
- Platt, J. R. Percentile reinforcement: paradigms for experimental analysis of response shaping. In G. H. Bower (Ed.) *The psychology of learning and motivation. Vol. 7*. New York: Academic Press, 1973. [118]
- Postman, L. Association and performance in the analysis of verbal learning. In T. R. Dixon and D. L. Horton (Eds.) *Verbal behavior and general behavior theory*. Englewood Cliffs, N. J.: Prentice-Hall, 1968. [286]
- Postman, L., and Keppel, G. (Eds.) *Verbal learning and memory*. Baltimore, Md.: Penguin Books, 1969. [277]
- Powers, R. B. Clock-delivered reinforcers in conjunctive and interlocking schedules. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 579-586. [185]
- Premack, D. Toward empirical behavior laws: I. Positive reinforcement. *Psychological Review*, 1959, 66, 219-233. [77, 226]
- Premack, D. Reversibility of the reinforcement relation. *Science*, 1962, 136, 255-257. [77]
- Premack, D. Reinforcement theory. In D. Levine (Ed.) *Nebraska symposium on motivation*. Lincoln: University of Nebraska Press, 1965. [77]
- Premack, D. A functional analysis of language. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 107-125. [248-249, 274]
- Premack, D. Catching up with common sense or two sides of a generalization: reinforcement and punishment. In R. Glaser, (Ed.) *The nature of reinforcement*. New York: Academic Press, 1971. [77-78, 99]

- Provine, R. R. Development of function in nerve nets. In J. Fentress (Ed.) *Simpler networks and behavior*. Sunderland, Mass.: Sinauer, 1976. [56]
- Pryor, K. W., Haag, R., and O'Reilly, J. The creative porpoise: training for novel behavior. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 653-661. [137]
- Prytulak, L. S. Natural language mediation. *Cognitive Psychology*, 1971, 2, 1-56. [328]
- Rachlin, H. The effect of shock intensity on concurrent and single-key responding in concurrent-chain schedules. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 87-93. [101, 196]
- Rachlin, H. On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 249-251. [193]
- Rachlin, H. Contrast and matching. *Psychological Review*, 1973, 80, 217-224. [188]
- Rachlin, H. Self-control. *Behaviorism*, 1974, 2, 94-107. [196, 350]
- Rachlin, H., and Baum, W. M. Effects of alternative reinforcement: does the source matter? *Journal of the Experimental Analysis of Behavior*, 1972, 18, 231-241. [194]
- Rachlin, H., and Frankel, M. Choice, rate of response, and rate of gambling. *Journal of Experimental Psychology*, 1969, 80, 444-449. [180]
- Rachlin, H., and Green, L. Commitment, choice and self-control. *Journal of the Experimental Analysis of Behavior*, 1972, 17, 15-22. [196]
- Ratner, S. C. Habituation: research and theory. In J. H. Reynierse (Ed.) *Current issues in animal learning*. Lincoln: University of Nebraska Press, 1970. [52]
- Reber, A. S. Locating clicks in sentences: Left, center, and right. *Perception and Psychophysics*, 1973, 13, 133-138. [262]
- Reber, A. S. Implicit learning of synthetic languages: the role of instructional set. *Journal of Experimental Psychology: Human Learning and Memory*, 1976, 2, 88-94. [247]
- Reber, A. S., and Anderson, J. R. The perception of clicks in linguistic and non-linguistic messages. *Perception and Psychophysics*, 1970, 8, 81-89. [262]
- Reid, L. S. The development of noncontinuity behavior through continuity learning. *Journal of Experimental Psychology*, 1953, 46, 107-112. [303]
- Reid, R. L. The role of the reinforcer as a stimulus. *British Journal of Psychology*, 1958, 49, 202-209. [81]
- Rescorla, R. A. Pavlovian conditioning and its proper control procedures. *Psychological Review*, 1967, 74, 71-80. [205]
- Rescorla, R. A. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, 1968, 66, 1-5. [219, 221]
- Rescorla, R. A., and Skucy, J. C. Effect of response-independent reinforcers during extinction. *Journal of Comparative and Physiological Psychology*, 1969, 67, 381-389. [85]
- Rescorla, R. A., and Solomon, R. L. Two process learning theory: relationships between Pavlovian conditioning and instrumental learning. *Psychological Review*, 1967, 74, 151-182. [215-217]
- Rescorla, R. A., and Wagner, A. R. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy (Eds.) *Classical conditioning II: current research and theory*. New York: Appleton-Century-Crofts, 1972. [210]

- Restle, F. Discrimination of cues in mazes: A resolution of the "place-versus-response" question. *Psychological Review*, 1957, 64, 217-228. [146]
- Restle, F. The selection of strategies in cue learning. *Psychological Review*, 1962, 69, 329-343. [346]
- Revusky, S. H. Mathematical analysis of the durations of reinforced interresponse times during variable interval reinforcement. *Psychometrika*, 1962, 27, 307-314. [177]
- Revusky, S. H., and Garcia, J. Learned associations over long delays. In G. H. Bower (Ed.), *The psychology of learning and motivation*. Volume 4. New York: Academic Press, 1970. [160-162]
- Reynolds, B. The acquisition of a trace conditioned response as a function of the stimulus trace. *Journal of Experimental Psychology*, 1945, 35, 15-30. [202]
- Reynolds, G. S. An analysis of interactions in a multiple schedule. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 107-117. (a) [188]
- Reynolds, G. S. Attention in the pigeon. *Journal of the Experimental Analysis of Behavior*. 1961, 4, 203-208. (b) [144-145]
- Reynolds, G. S. Behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1961, 4, 57-71. (c) [188]
- Riess, B. F. Genetic changes in semantic conditioning. *Journal of Experimental Psychology*, 1946, 36, 143-152. [248]
- Rilling, M. Stimulus control and inhibitory processes. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. New York: Prentice-Hall, 1977. [149]
- Risley, T. R., and Hart, B. Developing correspondence between the nonverbal and verbal behavior of preschool children. *Journal of Applied Behavior Analysis*, 1968, 1, 267-281. [246]
- Rizley, R. C., and Rescorla, R. A. Associations in second-order conditioning and sensory preconditioning. *Journal of Comparative and Physiological Psychology*, 1972, 81, 1-11. [209-210]
- Robinson, E. S., and Brown, M. A. Effect of serial position upon memorization. *American Journal of Psychology*, 1926, 37, 538-552. [283-284]
- Rock, I. The role of repetition in association learning. *American Journal of Psychology*, 1957, 70, 186-193. [287]
- Romney, A. K., and D'Andrade, R. G. Cognitive aspects of English kin terms. *American Anthropologist*, 1964, 66, 146-170. [271]
- Rosch, E. H. Natural categories. *Cognitive Psychology*, 1973, 4, 328-350. [159, 271]
- Rosenfeld, H. M., and Baer, D. M. Unbiased and unnoticed verbal conditioning: the double agent robot procedure. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 99-105. [244]
- Rozin, P., and Kalat, J. W. Specific hungers and poison avoidance as adaptive specializations in learning. *Psychological Review*, 1971, 78, 459-486. [160]
- Rumbaugh, D. M., Gill, T. V., von Glasersfeld, E., Warner, H., and Pisani, P. Conversations with a chimpanzee in a computer-controlled environment. *Biological Psychiatry*, 1975, 10, 627-641. [274]
- Rundus, D. Analysis of rehearsal processes in free recall. *Journal of Experimental Psychology*, 1971, 89, 63-77. [291]
- Rundus, D., and Atkinson, R. C. Rehearsal processes in free recall: a procedure for direct observation. *Journal of Verbal Learning and Verbal Behavior*, 1970, 9, 99-105. [291]
- Ryle, G. *The concept of mind*. New York: Barnes and Noble, 1949. [4]

- Sachs, J. S. Recognition memory for syntactic and semantic aspects of connected discourse. *Perception and Psychophysics*, 1967, 2, 437-442. [268-269, 335]
- Sakitt, B. Iconic memory. *Psychological Review*, 1976, 83, 257-276. [314]
- Salapatek, P., and Kessen, W. Visual scanning of triangles by the human newborn. *Journal of Experimental Child Psychology*, 1966, 3, 155-167. [91]
- Saugstad, P., and Raaheim, K. Problem solving, past experience and availability of functions. *British Journal of Psychology*, 1960, 51, 97-104. [350]
- Savin, H., and Perchonock, E. Grammatical structure and the immediate recall of English sentences. *Journal of Verbal Learning and Verbal Behavior*, 1965, 4, 348-353. [265]
- Schane, S. A., Tranel, B., and Lane, H. On the psychological reality of a natural rule of syllable structure. *Cognition*, 1975, 3, 351-358. [261]
- Schick, K. Operants. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 413-423. [245]
- Schlosberg, H. The relationship between success and the laws of conditioning. *Psychological Review*, 1937, 44, 379-394. [204, 207]
- Schneider, A. M. Control of memory by spreading cortical depression: a case for stimulus control. *Psychological Review*, 1967, 74, 201-215. [333]
- Schneider, B. A. A two-state analysis of fixed-interval responding in the pigeon. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 677-687. [181]
- Schoenfeld, W. N. An experimental approach to anxiety, escape and avoidance behavior. In P. H. Hoch (Ed.), *Anxiety*. New York: Grune and Stratton, 1950. [111]
- Schoenfeld, W. N. Some old work for modern conditioning theory. *Conditional Reflex*, 1966, 1, 219-223. [56]
- Schoenfeld, W. N. "Avoidance" in behavior theory. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 669-674. [111]
- Schoenfeld, W. N., Antonitis, J. J., and Bersh, P. J. Unconditioned response rate of the white rat in a bar-pressing apparatus. *Journal of Comparative and Physiological Psychology*, 1950, 43, 41-48. [82]
- Schoenfeld, W. N., and Cole, B. K. *Stimulus schedules: the t - τ systems*. New York: Harper and Row, 1972. [168, 185]
- Schoenfeld, W. N., Cumming, W. W., and Hearst, E. On the classification of reinforcement schedules. *Proceedings of the National Academy of Sciences*, 1956, 42, 563-570. [183]
- Schoenfeld, W. N., Harris, A. H., and Farmer, J. Conditioning response variability. *Psychological Reports*, 1966, 19, 551-557. [137]
- Scholz, K. W., and Potts, G. R. Cognitive processing of linear orderings. *Journal of Experimental Psychology*, 1974, 102, 323-326. [341]
- Schroeder, S. R., and Holland, J. G. Operant control of eye movements. *Journal of Applied Behavior Analysis*, 1968, 1, 161-166. [91]
- Schwartz, B. On going back to nature: a review of Seligman and Hager's *Biological boundaries of learning*. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 183-198. [160]
- Schwartz, B., and Gamzu, E. Pavlovian control of operant behavior. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [213]
- Schwartz, B., Hamilton, B., and Silberberg, A. Behavioral contrast in the pigeon: a study of the duration of key pecking maintained on multiple schedules of

- reinforcement. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 199-206. [190]
- Schwartz, B., and Williams, D. R. Two different kinds of key peck in the pigeon: some properties of responses maintained by negative and positive response-reinforcer contingencies. *Journal of the Experimental Analysis of Behavior*, 1972, 18, 201-216. [214]
- Sechenov, I. M. *Reflexes of the brain (1863)*. (tr. by S. Belsky) Cambridge, Mass.: The M.I.T. Press, 1965. [41, 55-56, 57-58]
- Segal, E. F. Induction and the provenance of operants. In R. M. Gilbert and J. R. Millenson (Eds.) *Reinforcement: behavioral analyses*. New York: Academic Press, 1972. [56, 358]
- Segal, E. F. Psycholinguistics discovers the operant: a review of Roger Brown's *A first language: the early stages*. *Journal of the Experimental Analysis of Behavior*, 1975, 23, 149-158. [274]
- Segal, E. F. Toward a coherent psychology of language. In W. K. Honig and J. E. R. Staddon (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [231]
- Seligman, M. E. P. On the generality of the laws of learning. *Psychological Review*, 1970, 77, 406-418. [110, 160]
- Seligman, M. E. P., and Hager, J. L. *Biological boundaries of learning*. New York: Appleton-Century-Crofts, 1972. [160]
- Sheffield, F. D. Relation between classical conditioning and instrumental learning. In W. F. Prokasy (Ed.) *Classical conditioning*. New York: Appleton-Century-Crofts, 1965. [113, 207-208]
- Shepard, R. N. Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior*, 1967, 6, 156-163. [295]
- Shepard, R. N., and Metzler, J. Mental rotation of three-dimensional objects. *Science*, 1971, 171, 701-703. [340-341]
- Sherrington, C. *The integrative action of the nervous system*. New York: Scribner's, 1906. [40]
- Shiffrin, R. M., and Atkinson, R. C. Storage and retrieval processes in long-term memory. *Psychological Review*, 1969, 76, 179-193. [314, 317, 325]
- Shimp, C. P. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 443-455. [193]
- Shimp, C. P. Synthetic variable-interval schedules of reinforcement. *Journal of the Experimental Analyses of Behavior*, 1973, 19, 311-330. [175]
- Shimp, C. P. Organization in memory and behavior. *Journal of the Experimental Analysis of Behavior*, 1976, 26, 113-130. [183, 335]
- Sidman, M. A note on functional relations obtained from group data. *Psychological Bulletin*, 1952, 49, 263-269. [67]
- Sidman, M. Two temporal parameters in the maintenance of avoidance behavior by the white rat. *Journal of Comparative and Physiological Psychology*, 1953, 46, 253-261. [108-109]
- Sidman, M. Delayed punishment effects mediated by competing behavior. *Journal of Comparative and Physiological Psychology*, 1954, 47, 145-147. [109]
- Sidman, M. *Tactics of scientific research*. New York: Basic Books, 1960. [86]
- Sidman, M. The behavioral analysis of aphasia. *Journal of Psychiatric Research*, 1971, 8, 413-422. [235]
- Sidman, M., and Boren, J. J. The relative aversiveness of warning signal and shock

- in an avoidance situation. *Journal of Abnormal and Social Psychology*, 1957, 55, 339-344. [192]
- Sidman, M., Herrnstein, R. J., and Conrad, D. G. Maintenance of avoidance behavior by unavoidable shocks. *Journal of Comparative and Physiological Psychology*, 1957, 50, 553-557. [217]
- Sidman, M., and Rosenberger, P. B. Several methods for teaching serial position sequences to monkeys. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 467-478. [154, 164]
- Sidman, M., and Stoddard, L. T. The effectiveness of fading in programming a simultaneous form discrimination for retarded children. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 3-15. [154]
- Silverman, P. J. Chained and tandem fixed-interval schedules of punishment. *Journal of the Experimental Analysis of Behavior*, 1971, 16, 1-13. [192]
- Sizemore, O. J., and Lattal, K. A. Dependency, temporal contiguity, and response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1977, 27, 119-125. [178-179]
- Skinner, B. F. On the conditions of elicitation of certain eating reflexes. *Proceedings of the National Academy of Sciences*, 1930, 16, 433-438. [68]
- Skinner, B. F. The concept of the reflex in the description of behavior. *Journal of General Psychology*, 1931, 5, 427-458. [22, 40]
- Skinner, B. F. The rate of establishment of a discrimination. *Journal of General Psychology*, 1933, 9, 302-350. [32]
- Skinner, B. F. A discrimination without previous conditioning. *Proceedings of the National Academy of Sciences*, 1934, 20, 532-536. (a) [32]
- Skinner, B. F. The extinction of chained reflexes. *Proceedings of the National Academy of Sciences*, 1934, 20, 234-237. (b) [134]
- Skinner, B. F. Two types of conditioned reflex and a pseudotype. *Journal of General Psychology*, 1935, 12, 66-77. [204, 207]
- Skinner, B. F. *The behavior of organisms*. New York: Appleton-Century-Crofts, 1938. [32-33, 41, 68, 80, 196]
- Skinner, B. F. "Superstition" in the pigeon. *Journal of Experimental Psychology*, 1948, 38, 168-172. [48-49, 84]
- Skinner, B. F. Are theories of learning necessary? *Psychological Review*, 1950, 57, 193-216. [68, 156]
- Skinner, B. F. *Science and human behavior*. New York: Macmillan, 1953. [88, 115, 338]
- Skinner, B. F. A case history in scientific method. *American Psychologist*, 1956, 11, 221-233. [68, 196]
- Skinner, B. F. *Verbal behavior*. New York: Appleton-Century-Crofts, 1957. [231, 242-243]
- Skinner, B. F. Diagramming schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 67-68. [185]
- Skinner, B. F. John Broadus Watson, behaviorist. *Science*, 1959, 129, 197-198. [340]
- Skinner, B. F. Behaviorism at fifty. *Science*, 1963, 140, 951-958. [326-327]
- Skinner, B. F. An operant analysis of problem-solving. In B. Kleinmuntz (Ed.) *Problem solving*. New York: Wiley, 1966. (a) [350]
- Skinner, B. F. The phylogeny and ontogeny of behavior. *Science*, 1966, 153, 1205-1213. (b) [358]

- Skinner, B. F. A lecture on "having" a poem. In B. F. Skinner, *Cumulative record (Third edition)*. New York: Appleton-Century-Crofts, 1972. [232]
- Skinner, B. F. The shaping of phylogenic behavior. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 117-120. [358]
- Skinner, B. F. *Particulars of my life*. New York: Alfred A. Knopf, 1976. [327]
- Slamecka, N. J. Recall and recognition in list-discrimination tasks as a function of the number of alternatives. *Journal of Experimental Psychology*, 1967, 74, 187-192. [296]
- Slamecka, N. J. Testing for associative storage in free recall. *Journal of Experimental Psychology*, 1969, 81, 557-560. [333]
- Slobin, D. I. *Psycholinguistics*. Glenview, Ill.: Scott, Foresman, 1971. [275]
- Small, W. S. Experimental studies of the mental processes of the rat. *American Journal of Psychology*, 1899-1900, 11, 1-89. [66-67]
- Smith, B. H. *Poetic closure*. Chicago: Chicago University Press, 1968. [232]
- Smith, J. B. Effects of response rate, reinforcement frequency, and the duration of a stimulus preceding response-independent food. *Journal of the Experimental Analysis of Behavior*, 1974, 21, 215-221. [219]
- Smith, K. Conditioning as an artifact. *Psychological Review*, 1954, 61, 217-225. [204, 207]
- Smith, K. R. *Behavior and conscious experience. A conceptual analysis*. Athens, Ohio: Ohio University Press, 1969. [4]
- Smith, W. G. The relation of attention to memory. *Mind*, 1895, 4, 47-73. [316]
- Snapper, A. G., Knapp, J. Z., and Kushner, H. K. Mathematical description of schedules of reinforcement. In W. N. Schoenfeld (Ed.) *The theory of reinforcement schedules*. New York: Appleton-Century-Crofts, 1970. [185]
- Solomon, P., Kubzansky, P. E., Leiderman, P. H., Mendelson, J. H., Trumbull, R., and Wexler, D. (Eds.) *Sensory deprivation*. Cambridge: Harvard University Press, 1961. [25]
- Solomon, R. L. Punishment. *American Psychologist*, 1964, 19, 239-253. [97]
- Solomon, R. L., and Corbit, J. D. An opponent-process theory of motivation: I. Temporal dynamics of affect. *Psychological Review*, 1974, 81, 119-145. [52, 61]
- Spalding, D. Instinct with original observations on young animals. *MacMillan's Magazine*, 1873, 27, 282-293. [57]
- Spence, K. W. The differential response in animals to stimuli varying within a single dimension. *Psychological Review*, 1937, 44, 430-444. [149, 304]
- Sperling, G. The information available in brief visual presentations. *Psychological Monographs*, 1960, 74 (11, Whole No. 498). [314-315]
- Sperling, G. A model for visual memory tasks. *Human Factors*, 1963, 5, 19-31. [314]
- Sperling, G. Successive approximations to a model for short term memory. *Acta Psychologica*, 1967, 27, 285-292. [314]
- Sperling, G., and Reeves, A. Reaction time of an unobservable response. Paper presented at the Psychonomic Society, St. Louis, Mo., November, 1976. [339]
- Staddon, J. E. R. Some properties of spaced responding in pigeons. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 19-27. [130]
- Staddon, J. E. R. Inhibition and the operant. A review of *Sensory inhibition by G. von Békésy, and Mach Bands: quantitative studies on neural networks in*

- the retina, by Floyd Ratliff. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 481-487. [188]
- Staddon, J. E. R., and Simmelhag, V. L. The "Superstition" experiment: a re-examination of its implications for the principles of adaptive behavior. *Psychological Review*, 1971, 78, 3-43. [49-51, 84, 358]
- Stein, L., Sidman, M., and Brady, J. V. Some effects of two temporal variables on conditioned suppression. *Journal of the Experimental Analysis of Behavior*, 1958, 1, 153-162. [219]
- Sternberg, S. High-speed scanning in human memory. *Science*, 1966, 153, 652-654. [343-345]
- Sternberg, S. Two operations in character recognition; some evidence from reaction time measurements. *Perception and Psychophysics*, 1967, 2, 45-53. [342]
- Sternberg, S. Memory-scanning; mental processes revealed by reaction time experiments. *American Scientist*, 1969, 57, 421-457. [342-345]
- Stewart, G. R. *Names on the land*. New York: Oxford, 1975. [241]
- Stratton, G. M. Vision without inversion of the retinal image. *Psychological Review*, 1897, 4, 341-360; 463-481. [89-90]
- Stratton, G. M. *Theophrastus and the Greek physiological psychology before Aristotle*. New York: Macmillan, 1917. [326]
- Stroop, J. R. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, 1935, 18, 643-662. [249]
- Stubbs, D. A., and Cohen, S. L. Second-order schedules: comparison of different procedures for scheduling paired and nonpaired brief stimuli. *Journal of the Experimental Analysis of Behavior*, 1972, 18, 403-413. [192]
- Taylor, A. B., and Irion, A. L. Continuity hypothesis and transfer of training in paired-associate learning. *Journal of Experimental Psychology*, 1964, 68, 573-577. [288]
- Terrace, H. S. Discrimination learning with and without "errors." *Journal of the Experimental Analysis of Behavior*, 1963, 6, 1-27. (a) [154]
- Terrace, H. S. Errorless transfer of a discrimination across two continua. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 223-232. (b) [154]
- Terrace, H. S. Stimulus control. In W. K. Honig (Ed.) *Operant behavior: Areas of research and application*. New York: Appleton-Century-Crofts, 1966. [188]
- Terrace, H. S. Discrimination learning, the peak shift, and behavioral contrast. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 727-741. [188]
- Terrace, H. S. Evidence of the innate basis of the hue dimension in the duckling. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 79-87. [151]
- Thistlethwaite, D. A critical review of latent learning and related experiments. *Psychological Bulletin*, 1951, 48, 97-129. [86]
- Thomas, D. R., and Switalski, R. W. Comparison of stimulus generalization following variable-ratio and variable-interval training. *Journal of Experimental Psychology*, 1966, 71, 236-240. [147]
- Thomas, J. C., Jr. An analysis of behavior in the hobbits-orcs problem. *Cognitive Psychology*, 1974, 6, 257-269. [348-349]
- Thorndike, E. L. Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monograph Supplements*, 2 (No. 4), 1898, pp. 109. [29-30, 64-65]

- Thorndike, E. L. *The psychology of learning*. New York: Teacher's College, 1913. [56]
- Thorndike, E. L., and Woodworth, R. S. The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review*, 1901, 8, 247-261. [298]
- Tighe, T. J., Glick, J., and Cole, M. Subproblem analysis of discrimination shift learning. *Psychonomic Science*, 1971, 24, 159-160. [304]
- Tinbergen, N. *The Herring Gull's world*. (Revised edition) New York: Basic Books, 1960. [26-28]
- Tinbergen, N., and Perdeck, A. C. On the stimulus situation releasing the begging response in the newly hatched Herring Gull chick (*Larus a. argentatus* Pontopp.). *Behaviour*, 1950, 3, 1-38. [28]
- Titchener, E. B. The postulates of a structural psychology. *Philosophical Review*, 1898, 7, 449-465. [231, 356]
- Tolman, E. C. Cognitive maps in rats and men. *Psychological Review*, 1948, 55, 189-208. [86, 146]
- Tolman, E. C., and Honzik, C. H. Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology*, 1930, 4, 257-275. [86-87]
- Touchette, P. E. Tilted lines as complex stimuli. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 211-214. [153]
- Townsend, J. T. A note on the identifiability of parallel and serial processes. *Perception and Psychophysics*, 1971, 10, 161-163. [345]
- Tracy, W. K. Wavelength generalization and preference in monochromatically reared ducklings. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 163-178. [151]
- Trehub, S. E., and Chang, H.-W. Speech as reinforcing stimulation for infants. *Developmental Psychology*, 1977, 13, 170-171. [232]
- Tulving, E. Subjective organization in free recall of "unrelated" words. *Psychological Review*, 1962, 69, 344-354. [272, 292]
- Tulving, E. Retrograde amnesia in free recall. *Science*, 1969, 164, 88-90. [322]
- Tulving, E. Episodic and semantic memory. In E. Tulving and W. Donaldson (Eds.) *Organization of memory*. New York: Academic Press, 1972. [334]
- Tulving, E. Cue-dependent forgetting. *American Scientist*, 1974, 62, 74-82. [297, 332]
- Tulving, E., and Madigan, S. A. Memory and verbal learning. *Annual Review of Psychology*, 1970, 21, 437-484. [277, 320, 330-331]
- Tulving, E., and Pearlstone, Z. Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 1966, 5, 381-391. [331]
- Tulving, E., and Psotka, J. Retroactive inhibition in free recall: inaccessibility of information available in the memory store. *Journal of Experimental Psychology*, 1971, 87, 1-8. [331-332]
- Tweney, R. D. Review of *Analogy and association in linguistics and psychology* by Erwin A. Esper. *Historiographia Linguistica*, 1974, 1, 385-398. [275]
- Twitmyer, E. B. A study of the knee jerk (1902). *Journal of Experimental Psychology*, 1974, 103, 1047-1066. [205]

- Underwood, B. J. Interference and forgetting. *Psychological Review*, 1957, 64, 49-60. [322-323]
- Underwood, B. J. Ten years of massed practice on distributed practice. *Psychological Review*, 1961, 68, 229-247. [283]
- Underwood, B. J. Degree of learning and measurement of forgetting. *Journal of Verbal Learning and Verbal Behavior*, 1964, 3, 112-129. [321]
- Underwood, B. J., and Freund, J. S. Errors in recognition learning and retention. *Journal of Experimental Psychology*, 1968, 78, 55-63. [296]
- Underwood, B. J., and Goad, D. Studies of distributed practice: I. The influence of intra-list similarity in serial learning. *Journal of Experimental Psychology*, 1951, 42, 125-134. [283]
- Underwood, B. J., and Postman, L. Extraexperimental sources of interference in forgetting. *Psychological Review*, 1960, 67, 73-95. [323]
- Underwood, B. J., Rehaula, R., and Keppel, G. Item-selection in paired-associate learning. *American Journal of Psychology*, 1962, 75, 353-371. [288]
- Underwood, B. J., and Schulz, R. W. Meaningfulness and verbal learning. Philadelphia: Lippincott, 1960. [283, 286]
- van Lawick-Goodall, J. The behavior of free-living chimpanzees in the Gombe Stream Reserve. *Animal Behaviour Monographs*, 1968, 1, 161-311. [26]
- Verhave, T. Contributions to the history of psychology: IV. Joseph Buchanan (1785-1829) and the "law of exercise" (1812). *Psychological Reports*, 1967, 20, 127-133. [56]
- von Restorff, H. Über die Wirkung von Bereichsbildungen im Spurenfeld. *Psychologische Forschung*, 1933, 18, 299-342. [290]
- Voss, S. C., and Homzie, M. J. Choice as a value. *Psychological Reports*, 1970, 26, 912-914. [13]
- Wagner, A. R., Thomas, E., and Norton, T. Conditioning with electrical stimulation of motor cortex: evidence of a possible source of motivation. *Journal of Comparative and Physiological Psychology*, 1967, 64, 191-199. [207]
- Walters, G. C., and Glazer, R. D. Punishment of instinctive behavior in the mongolian gerbil. *Journal of Comparative and Physiological Psychology*, 1971, 75, 331-340. [101]
- Wang, M. D. Frequency effects in the abstraction of linguistic ideas. *Bulletin of the Psychonomic Society*, 1977, 9, 303-306. [269]
- Wason, P. C. The contexts of plausible denial. *Journal of Verbal Learning and Verbal Behavior*, 1965, 4, 7-11. [333]
- Wasserman, E., Franklin, S., and Hearst, E. Pavlovian appetitive contingencies and approach vs. withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology*, 1974, 86, 616-627. [214]
- Watkins, M. J., and Tulving, E. Episodic memory: when recognition fails. *Journal of Experimental Psychology: General*, 1975, 104, 5-29. [297]
- Watson, J. B. *Psychology from the standpoint of a behaviorist*. Philadelphia: Lippincott, 1919. [40]
- Watson, J. B., and Rayner, R. Conditioned emotional reactions. *Journal of Experimental Psychology*, 1920, 3, 1-14. [206]
- Waugh, N. C. Retention as an active process. *Journal of Verbal Learning and Verbal Behavior*, 1972, 11, 129-140. [324]
- Waugh, N. C., and Norman, D. A. Primary memory. *Psychological Review*, 1965, 72, 89-104. [314, 316, 319]

- Weiner, H. Controlling human fixed-interval performance. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 349-373. [180]
- Weisberg, R., and Suls, J. M. An information-processing model of Duncker's candle problem. *Cognitive Psychology*, 1973, 4, 255-276. [350]
- Weiss, B., and Laties, V. G. Behavioral thermoregulation. *Science*, 1961, 133, 1338-1344. [107]
- Wendt, G. R. An interpretation of inhibition of conditioned reflexes as competition between reaction systems. *Psychological Review*, 1936, 43, 258-281. [194]
- Wertheimer, M. *Productive thinking*. New York: Harper and Row, 1959. [348]
- Whorf, B. L. *Language, thought, and reality*. Cambridge, Mass.: The M.I.T. Press, 1956. [241]
- Wickelgren, W. A. Context-sensitive coding, associative memory, and serial order in (speech) behavior. *Psychological Review*, 1969, 76, 1-15. [294]
- Wickens, D. D. Encoding categories of words; an empirical approach to meaning. *Psychological Review*, 1970, 77, 1-15. [320]
- Wildemann, D. G., and Holland, J. G. Control of a continuous response dimension by a continuous stimulus dimension. *Journal of the Experimental Analysis of Behavior*, 1972, 18, 419-434. [159]
- Wilkins, A. T. Conjoint frequency, category size, and categorization time. *Journal of Verbal Learning and Verbal Behavior*, 1971, 10, 382-385. [272]
- Williams, B. A. The blocking of reinforcement control. *Journal of the Experimental Analysis of Behavior*, 1975, 24, 215-225. [194]
- Williams, B. A. The effects of unsignalled delayed reinforcement. *Journal of the Experimental Analysis of Behavior*, 1976, 26, 441-449. [179]
- Williams, D. R., and Williams, H. Auto-maintenance in the pigeon: sustained pecking despite contingent non-reinforcement. *Journal of the Experimental Analysis of Behavior*, 1969, 12, 511-520. [213]
- Wilson, M. P., and Keller, F. S. On the selective reinforcement of spaced responses. *Journal of Comparative and Physiological Psychology*, 1953, 46, 190-193. [130]
- Wollen, K. A. Effects of maximizing availability and minimizing rehearsal upon associative symmetry in two modalities. *Journal of Experimental Psychology*, 1968, 77, 626-630. [288]
- Wollen, K. A., Weber, A., and Lowry, D. Bizarreness versus interaction of mental images as determinants of learning. *Cognitive Psychology*, 1972, 3, 518-523. [311]
- Woodward, A., Jr., and Murdock, B. B., Jr. Positional and sequential probes in serial learning. *Canadian Journal of Psychology*, 1968, 22, 131-138. [283]
- Woodworth, R. S. *Psychology* (Revised Edition). New York: Holt, 1921. [307, 338]
- Woodworth, R. S. *Experimental Psychology*. New York: Holt, 1938. [299]
- Woodworth, R. S., and Schlosberg, H. *Experimental Psychology*, New York: Holt, Rinehart and Winston, 1954. [65]
- Wright, A. A., and Cumming, W. W. Color-naming functions for the pigeon. *Journal of the Experimental Analysis of Behavior*, 1971, 15, 7-17. [156]
- Wundt, W. *Die Sprache*. Leipzig: Engelmann, 1900. [259]
- Yates, F. A. *The art of memory*. Chicago: University of Chicago Press, 1966. [311-312]
- Yerkes, R. M. *The dancing mouse*. New York: Macmillan, 1907. [33]

- Yerkes, R. M., and Watson, J. B. Methods of studying vision in animals. *Behavior Monographs*, 1911, 1 (No. 2, iv) pp. 90. [33]
- Yntema, D. B., and Trask, F. P. Recall as a search process. *Journal of Verbal Learning and Verbal Behavior*, 1963, 2, 65-74. [333]
- Young, F. A. Studies of pupillary conditioning. *Journal of Experimental Psychology*, 1958, 55, 97-110. [206]
- Zangwill, O. L. Remembering revisited. *Quarterly Journal of Experimental Psychology*, 1972, 24, 123-138. [310]
- Zeigarnik, B. Das Behalten erledigter und unerledigter Handlungen. *Psychologische Forschung*, 1927, 9, 1-85. [325]
- Zeiler, M. D. Fixed and variable schedules of response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, 11, 405-414. [178]
- Zeiler, M. D. Schedules of reinforcement: the controlling variables. In W. K. Honig and J. E. R. Staddon, (Eds.) *Handbook of operant behavior*. Englewood Cliffs, N. J.: Prentice-Hall, 1977. [181, 186]
- Zelazo, P. R., Zelazo, N. A., and Kolb, S. Newborn walking. *Science*, 1972, 177, 1058-1060. (a) [58]
- Zelazo, P. R., Zelazo, N. A., and Kolb, S. "Walking" in the newborn. *Science*, 1972, 176, 314-315. (b) [58]
- Zener, K., and McCurdy, H. G. Analysis of motivation factors in conditioned behavior. I. Differential effect of change in hunger upon conditioned, unconditioned and spontaneous salivary secretion. *Journal of Psychology*, 1939, 8, 321-350. [57, 121]
- Zimmerman, J., Hanford, P. H., and Brown, W. Effects of conditioned reinforcement frequency in an intermittent free-feeding situation. *Journal of the Experimental Analysis of Behavior*, 1967, 10, 331-340. [190]
- Zuriff, G. E. A comparison of variable-ratio and variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1970, 13, 369-374. [175]

Acknowledgments

Thanks go to the sources listed below, as well as to the authors involved, for granting permission to use material from copyrighted works. Full citations appear in the References section.

Fig. 2-2, from Tinbergen and Perdeck (1950), published with permission of E. J. Brill Publishers. Quotation on pages 32-33 from Skinner (1938), copyright 1938 by Appleton-Century-Crofts; reprinted by permission of Prentice-Hall, Inc. Fig. 2-4, from Lashley (1930), by permission of The Journal Press.

Fig. 3-3, from Staddon and Simmelhag (1971), copyright 1971 by the American Psychological Association; adapted by permission.

Fig. 4-10, from Tolman and Honzik (1930), published in 1930 by The Regents of the University of California; reprinted by permission of the University of California Press. Fig. 4-11, from Held and Hein (1963), copyright 1963 by the American Psychological Association; reprinted by permission.

Fig. 5-3, from Camp, Raymond, and Church (1967), copyright 1967 by the American Psychological Association; adapted by permission. Fig. 5-4, from Holz and Azrin (1961), copyright 1961 by the Society for the Experimental Analysis of Behavior, Inc.; published with permission of the Society. Fig. 5-5, from Fowler and Trapold (1962), copyright 1962 by the American Psychological Association; adapted by permission. Fig. 5-7, from Sidman (1953b), copyright 1953 by the American Psychological Association; reprinted by permission.

Fig. 6-1, from Miller and Banuazizi (1968), copyright 1968 by the American Psychological Association; reprinted by permission. Fig. 6-4, from Hefferline and Kennan (1963), copyright 1963 by the Society for the Experimental Analysis of Behavior, Inc.; published with permission of the Society.

Fig. 7-1, from Herrick, Myers, and Korotkin (1959), copyright 1959 by the American Psychological Association; reprinted by permission. Fig. 7-2, from Reynolds (1961b), copyright 1961 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 7-3, from Jenkins and Harrison (1960), copyright 1960 by the American Psychological Association; reprinted by permission. Fig. 7-4, from Hanson (1959), copyright 1959 by the American Psychological Association; adapted by permission. Fig. 7-5, from Honig, Boneau, Burstein, and Pennypacker (1963), copyright 1963 by the American Psychological Association; reprinted by permission. Fig. 7-7, from Revusky and Garcia (1970), published with permission of Academic Press.

Fig. 8-2, from Catania and Reynolds (1968), copyright 1968 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 8-5, from Catania, Matthews, Silverman, and Yohalem (1977), copyright 1977 by the Society for the

Experimental Analysis of Behavior, Inc.; published with permission of the Society. Fig. 8-7, from Sizemore and Lattal (1977), copyright 1977 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 8-8, from Dews (1962), copyright 1962 by the Society for the Experimental Analysis of Behavior, Inc.; published with permission of the Society. Fig. 8-10, from Reynolds (1961c), copyright 1961 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 8-11, from Catania (1969), copyright 1969 by the Society for the Experimental Analysis of Behavior, Inc.; published with permission of the Society.

Fig. 9-7, from Geller (1960), copyright 1960 by the Society for the Experimental Analysis of Behavior, Inc.; published with permission of the Society. Fig. 9-8 from Blackman (1968a), copyright 1968 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 9-9, from Blackman (1968b), copyright 1968 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society. Fig. 9-10, from Rescorla (1968), copyright 1968 by the American Psychological Association; adapted by permission.

Fig. 10-3, from Premack (1970), copyright 1970 by the Society for the Experimental Analysis of Behavior, Inc.; adapted with permission of the Society.

Fig. 11-2, from Johnson (1965), adapted with permission of Academic Press. Fig. 11-4, from Collins and Quillian (1969), adapted with permission of Academic Press.

Fig. 12-4, from Rundus and Atkinson (1970), adapted with permission of Academic Press. Table on page 295, from Goldiamond (1966), copyright © 1966 by John Wiley & Sons, Inc.; adapted by permission of John Wiley & Sons, Inc.

Fig. 13-1, from Sperling (1960), copyright 1960 by the American Psychological Association; adapted by permission. Fig. 13-2, from Peterson and Peterson (1959), copyright 1959 by the American Psychological Association; adapted by permission. Fig. 13-3, from Melton (1963), published with permission of Academic Press. Fig. 13-5, from Underwood (1957), copyright 1957 by the American Psychological Association; reprinted by permission. Fig. 13-7, from Tulving (1974), published with permission of *American Scientist*.

Fig. 14-1, from Shepard and Metzler (1971), copyright 1971 by the American Association for the Advancement of Science. Fig. 14-2, from Sternberg (1969), by permission of the author. Fig. 14-3, from Sternberg (1966), copyright 1966 by the American Association for the Advancement of Science. Fig. 14-4, from Thomas (1964), adapted with permission of Academic Press.

Index

- Abbreviations, schedule, 184-185, 187
- Abstraction, 156, 159, 261
 - in tacting, 237-240, 347
- Acceleration, 71-72, 79
- Accessibility in memory, 308, 321, 326, 330
- Accommodation, 346
- Activity wheel, 14, 49, 77, 99
- Adaptation, 39, 52-55, 59, 61, 200
- Afterimage, 314
- Aggression, 82-83, 216
- Alcohol, 49
- Alphabet, 156, 235, 242, 303
- Alternative schedule, 185, 187
- Analysis of behavior, 7, 35, 103, 136, 183, 227, 230
- Anatomy, 12fn, 273, 355-356
- Anxiety, 199, 215-216
- Aphasia, 235
- Approximations, successive, 117-119, 153-155 (*see also* Fading, Shaping)
- Approximation to English, 257
- Arithmetic, 253
- Assertion, 251
- Assimilation, 346
- Associations, 63, 204, 282, 355, 359
 - backward, 284-285
 - in chains, 134-135
 - contingent, 294
 - free, 242-243
 - and memory, 321
 - in paired-associates learning, 286-289, 297, 303
 - remote, 284-285
 - R-S vs. S-S, 88
 - symmetry of, 288
 - word, 257, 270
- Attention:
 - in conditioning, 210
 - as dimensional control, 143-145, 148, 153-154, 278
 - as response, 338-340, 354
- Audience, 231, 237, 240, 244
- Autoclitics, 239, 248, 250-253, 268, 292
 - descriptive, 251
 - qualifying and quantifying, 251-252
 - relational, 252-253
- Automaintenance, 212-215
- Autonomic behavior, 47-48, 248
 - operant control of, 120-122, 207, 212
- Autoshaping, 212-215, 217, 354
 - and contrast, 190, 214
- Availability in memory, 308, 321, 326, 331
- Aversion, food, 160, 355
- Aversive control (*see* Avoidance, Escape, Punishment)
- Avoidance, 104-105, 108-113, 207, 226
 - continuous, 109
 - extinction, 112-113
 - passive, 98, 114
 - and preaversive stimuli, 217
 - Sidman, 108-111
 - signalled, 57, 226
 - two-process theories, 110-112, 220
- Awareness, 127
- Babinski reflex, 43-44
- Bait-shyness, 160-162
- Baseline, 20, 82, 170
 - in contrast, 189-190
 - in operant-respondent interaction, 215, 218-219
 - in punishment research, 95-96
- Behavior, 17-37
 - adjunctive, 39, 49, 59, 353
 - free-operant, 70
 - instrumental, 47
 - interim, 39, 50-51, 213, 353
 - operant, 47, 207
 - respondent, 47, 207
 - social, 6, 20
 - spontaneous, 56-57, 121

- Behavior (*cont.*)
- terminal, 39, 50-51, 213, 353
 - voluntary, 47, 120
 - (*see also* Autonomic behavior, Sequential behavior, Species-specific behavior, Verbal behavior)
- Behavioral contrast, 186-190 (*see also* Contrast)
- Behavior hierarchy, 18-21 (*see also* Hierarchy)
- Bilingualism, 288, 335
- Biology, 357-360
- Blocking, 210-211
- Bonds, S-R, 134, 135
- Braille, 231, 235
- Brain, 5-6
 - electrical stimulation of, 121-122
- Break-and-run responding, 173, 182
- Cat, 30-31, 64, 89
- Causation, 13, 51, 63
 - and contiguity, 177-180
- Chain, response, 85-86 (*see also* Sequential behavior)
- Chained schedule, 9, 186-187, 190-192
 - concurrent, 195-196
- Chamber, experimental (Skinner), 68-70
- Change in behavior (*see* Process)
- Changeover delay, 193
- Chick, 26-29, 30, 56-57
- Child, 115, 245-247, 274, 304, 346-347 (*see also* Infant)
- Chimpanzee, 25-26, 64, 192, 248-249, 274
- Choice, 13, 33, 78, 132
 - and concurrent schedules, 186, 192-196
- Chunking, 328
- Class, behavior, 227, 358-361 (*see also* Operant, Respondent)
 - and memory, 335
 - response, 23, 119-123
 - verbal, open and closed, 273
- Classical conditioning (*see* Conditioning)
- Clustering, 292, 297
- Cognition, 11, 253, 268, 277, 337-347, 358-360
- Cold, as aversive stimulus, 107-108, 225, 354
- Componential analysis, 271
- Comprehension, verbal, 245, 268
- Concept, 17, 273
 - formation, 159, 346, 354
 - natural, 159-160
 - in facting, 238
 - in verbal discrimination, 293
- Concurrent schedules, 183, 186-187, 192-196
 - chained, 195-196
- Conditional discrimination, 240, 293
- Conditional reflex, 32, 40, 42, 45, 134, 199-215
- Conditional response (*see* CR)
- Conditional stimulus (*see* CS)
- Conditioned aversive stimulus, 61
- Conditioned emotional reaction, 206
- Conditioned emotional response (CER), 215
- Conditioned reflex (*see* Conditional reflex)
- Conditioned reinforcement, 61, 186, 190-192
- Conditioned suppression, 215, 354
- Conditioning, 31-32, 63, 75, 199-221, 226, 360
 - backward, 203-204
 - and consequences, 205-208
 - defensive, 206-207
 - delay, 202-203
 - differential, 203-204
 - second-order, 208-210
 - semantic, 248
 - simultaneous, 202-203
 - not stimulus substitution, 204, 210, 220
 - temporal, 203-204
 - trace, 202-203
- Conjoint schedule, 187
- Conjunctive schedule, 185-187
- Consciousness, 167, 316
 - motor theory of, 234
- Consequences, 24, 29-37, 56-57, 60-115, 117, 277
 - in conditioning, 205-208
 - natural, 78, 246-247
 - in paradigm, 222-223
 - and stimulus control, 32-37
 - and verbal behavior, 235, 237, 244, 278
- Consequential operation (*see* Consequences)
- Conservation, 347
- Constituents, 256, 258-262, 267
- Constraints on learning (*see* Limits of learning)
- Contiguity, 63, 177-180, 282, 355
 - and conditioning, 204-205, 214, 219
- Contingency, 63, 82, 334, 360
 - in conditioning, 204-205, 214
 - defined, 82, 222, 312
 - and extinction, 82-83, 112-113
 - and instructions, 247
 - mnemonic for, 312
 - paradigm, 221-226
 - in problem-solving, 350-351
 - in stimulus-control, 142, 222-223
 - three-term, 22, 32-33, 165
- Continuity-discontinuity, 287
- Contrast, behavioral, 141, 186-190, 196, 360
 - and autoshaping, 190, 214
- Copying, not transcription, 233

- Correction procedure, 158
- Correlation:
 - in conditioning, 204-205, 215, 219-220
 - in schedules, 177-180
- CR, 201-203, 208-215
- Creativity, 330
- Critical features, 27-28, 156
- Critical period in imprinting, 59
- CS, 201-204, 208-215
- Cue-dependency, 308, 331-333
- Cumulative record, 70-73, 80, 81, 85
 - VI and VR, 169-170, 172-174
- Curiosity, 89
- CVC-trigram, 282, 287

- Defense reaction, 110
- Delay:
 - changeover, 193
 - conditioning, 202-203
 - of reinforcement, 105, 178-179, 181, 195
- Demand characteristics, 278
- Deprivation, 60-61, 77-78, 98, 199
 - and generalization, 147
 - sensory, 25
- Descartes, R., 40, 42
- Description, 15, 75, 91, 232, 329, 359
- Determinism, 13
- Development, 357, 359
 - language, 268, 274-275
 - stages of, 346-347
- Dictation (taking), 234-235, 236, 301
- Differential reinforcement, 118-127, 136-137, 155
 - of IRTs, 175-177
 - of latency, 127-129
 - of pauses, 179fn
 - schedules of, 129-131
 - in stimulus control, 141, 143
 - of successive approximations (*see* Shaping)
- Differential-reinforcement-of-rate schedules (*see* DRH, DRL, DRO, DRP)
- Differentiation, 118-127, 139, 142-143, 277, 360
 - and metamemory, 324
 - vocabulary, 155
- Dimensional control, 144-145, 147-149, 151
- Discriminated operant, 139-143, 155-160, 163-165, 358
- Discrimination, 32-37, 133, 139, 142-165, 226, 360
 - conditional, 240, 293
 - errorless, 154
 - and experience, 150-151
 - feature-positive, 153
 - and metamemory, 324
 - and multiple schedule, 186-190
 - reversal, 303
 - simultaneous, 35, 292-293
 - successive, 35
 - and transfer, 303
 - verbal (*see* Verbal discrimination)
 - vocabulary, 155
- Discriminative stimulus, 32, 141
 - and problem-solving, 350-351
 - punishment as, 102-103
 - reinforcer as, 83
- Distinctiveness, 145, 290-291
- Distractor items, 296
- Dog, 30-32, 43, 46-48, 57, 121, 200-212
- Double alternation, 132, 137
- DRH schedule, 130, 184, 218
- Drinking, 19, 49, 77, 99, 160
 - in autoshaping, 214
- DRL schedule, 130-131, 143, 167, 175-176, 183-185, 217-219
- DRO schedule, 131, 185
- DRP schedule, 131, 185, 218
- Duckling, 59-60, 151

- Ebbinghaus, H., 323
- Echoic behavior, 231-232, 236, 240, 242, 247, 301
- Echoic memory, 316
- Effect, law of, 64, 73, 91, 96
- Electrical stimulation of brain, 121-122
- Electric shock, 68, 99, 101-102
 - and conditioning, 206-207
 - inescapable and unavoidable, 164, 217
 - response-produced and response-independent, 100
- Electroconvulsive shock, 322, 333
- Elicitation, 24, 28, 38-61
 - and conditioning, 199-215
 - and contrast, 190
 - and escape, 106-107
 - and punishment, 100-102
 - and reinforcement, 76, 83-84
 - and stimulus control, 31-32, 140
- Emission, 33-36, 39, 46, 200
- Emotion, 52, 83, 199-200, 206, 215-217
 - anger, 216
 - fear, 206, 209-210, 215-216
 - joy, 217
 - language of, 239
- Encoding, 308, 327, 332
 - acoustic, 327
 - and depth of processing, 329, 335
 - by elaboration, 327-328
 - and imagery, 341
 - semantic, 328-330
 - by substitution, 327-328
- Enhancement, conditioned, 217

- Escape, 57, 93, 105-108, 353
 - and conditioning, 207
 - and elicited responding, 106-107
 - extinction, 112
- Ethics, 97, 98, 115
- Ethology, 28
- Etymology, 6, 240
- Event record, 44, 72
- Evolution, 110, 112, 357-361
- Excitation, 149, 152, 189
- Exercise, 39, 55-59, 353
- Explanation, 11, 22, 75, 148-149, 216, 340, 359
- Exploratory behavior, 89
- Extinction, 75, 79-85, 95, 98, 184
 - in conditioning, 209
 - and gradients, 147
 - and inhibition, 79-82, 153
 - in multiple schedule, 140-141
 - and negative reinforcement, 112-113
 - and recovery from punishment, 95, 98
 - and superstition, 83-85
 - after VI and VR, 172-173
- Eye-blink, 44
- Eye movement, 21, 64, 90-91
 - and attention, 339-340
- Facilitation, 52-55, 59
- Fading, 153-155, 165
 - and errorless discrimination, 154
 - and verbal discrimination, 294-295
- Fatigue, 81
- Feedback, 278
- Feeder, 68-69
- FI schedule, 144, 167, 180-183, 184-185, 218, 247
- Fish, tropical, 55
- Fixation, 303
- Fixed action pattern, 28
- Fixed-interval schedule (*see* FI schedule)
- Fixed-ratio schedule (*see* FR schedule)
- Fixed-time schedule (*see* FT schedule)
- Forgetting, 283, 307, 308, 320-323
 - instructed, 324-325
- Formal classes, verbal, 231-236, 301
- Freedom, 13, 196
- Free-operant behavior, 70
- Free recall, 280-281, 289-292, 296-297, 324
 - clustering, 297, 331
 - modified, 289
- Frequency, relative (*see* Probability)
- Frog, 53
- FR schedule, 8, 180-183, 184
- Frustration, 81, 188
- FT schedule, 180, 183, 184-185
- Function, 11-12, 23, 227, 337, 356-358,
 - of memory, 335
 - in stimulus control, 165
 - of verbal behavior, 230-253, 256, 258
- Function fixity, 350
- Fuzzy set, 159, 271
- Galvanic skin response, 52, 248
- Generalization, 139, 142-143, 144, 147-153, 165
 - and deprivation, 147
 - and schedule, 147
 - in verbal behavior, 241
 - vocabulary, 155
- Generalization gradient, 147-149, 151
- Genetics, 359
- Gerbil, 101
- Glandular secretion, 47, 120
- Gradient of stimulus control, 146-153, 155, 170
 - and deprivation, 147
 - difference, 149
 - excitatory, 149, 152
 - generalization, 147-149, 151
 - inhibitory, 149, 151-153
 - post-discrimination, 149-151, 153
 - and schedule, 147
- Grain, of cumulative record, 71-72
- Grammar, 137, 230, 255-268, 275
- Gull, 26-29, 56-57
- Habit, 17
- Habituation, 52-55
- Hallucination, 25, 338
- Hampton Court maze, 66
- Heart rate, 121-122
- Heat, as reinforcer, 107-108, 225, 354
- Helplessness, learned, 164
- Herring gull (*see* Gull)
- Hierarchical organization in language, 266-267, 356
 - semantic, 271-272, 335
- Hierarchy:
 - response, 18-21, 61, 78, 99, 117, 360
 - habit-family, 18
- Higher-order schedule, 183
- Hobbits-and-orcs problem, 348-349
- Houselight, 68, 69, 181
- Hume, David, 282
- Hypotheses, 133
- Imagery and imagining, 338-342, 354
 - eidetic, 341
 - in mnemonics, 328
- Imitation, 26, 164, 231
- Imprinting, 39, 59-60
- Incubation, 329
- Induction, 119-127, 139, 142-143, 155
- Infant, 21-25, 43-44, 55-58, 91, 206

- Inhibition, 17, 151
 - in conditioning, 211-212
 - in contrast, 189
 - in extinction, 79-82, 153
 - proactive, 302-303, 319 (*see also* Interference)
 - reflex, 46, 59
 - retroactive, 302-303, 323
- Inhibitory gradient, 149, 151-153
- Innate behavior, 13, 26-28, 110 (*see also* Species-specific behavior)
- Insight, 26, 29, 64, 337
- Instinct, 39, 57 (*see also* Species-specific behavior)
- Instructions, 245-247, 250
 - to forget, 324-325
 - and memory, 331
 - and verbal learning, 278, 284, 289, 291, 297
- Intelligence, 26, 29, 64, 229
- Interaction, operant-respondent, 215-221
- Interference, 139, 277, 283
 - and extinction, 81
 - and memory, 308, 321-323, 331
 - proactive, 302-303, 319-320, 322-323, 324, 333
 - retroactive, 302-303, 323
- Interlocking schedule, 185-187
- Interresponse time (*see* IRT)
- Interval schedule (*see* FI schedule, VI schedule)
- Intestinal motility, 121-122, 214
- Intraverbal behavior, 242-244, 247, 253, 257-258, 277
- Introspection, 338
- IRT, differential reinforcement of, 130-131, 143, 185
 - in VI and VR, 175-177
- Jumping stand (Lashley), 33-34
- Kernel, 264, 266
- Key, pigeon, 8, 69, 117-119
- Kitten, 89
- Knee jerk (*see* Reflex, patellar)
- Knowledge, 4, 12, 337, 357
 - of results, 278
- Köhler, W., 64, 351
- Language, 134-135, 229-275
 - development, 268, 274-275
 - universals, 273-274
- Latency, 41
 - differentiation, 127-129
- Latent learning, 86-88, 226, 353, 357-358
- Laughing gull (*see* Gull)
- Laws (*see* Effect, Exercise, Matching, Reflex)
- Learned helplessness, 164
- Learning:
 - all-or-none, 65, 85, 287
 - intentional *vs.* incidental, 278
 - latent, 86-88, 226, 353, 357-358
 - observational, 164
 - place- *vs.* response-, 146
 - sensory-motor, 88-91
 - set, 163-165, 299
 - verbal (*see* Paired-associates learning, Recall, Recognition, Serial learning, Verbal discrimination)
- Lexicon, 229, 255
- LH, 177, 184, 218
- Light, as aversive stimulus, 107
- Limited hold (*see* LH)
- Limits of learning, 135-136, 355
 - and language, 273
- Linguistics, 229
- Link, 10, 195
- Listening, 244, 246, 247-250
- LSD, 99
- Maintenance of behavior, 8, 196
- Mand, 245, 248, 251
- Masochism, 102
- Massed practice, 283
- Matching, in concurrent schedules, 192-194
- Matching law, 193
- Matching-to-sample, 156-159, 248
 - symbolic, 158
- Maximizing, 192-194
 - momentary, 193
- Maze, 31, 64-68, 86-88, 131-133, 139, 146, 287, 357-358
 - T-maze, 66-67, 132, 146, 287
- Meaning, 135, 230, 248-249, 255, 263, 268-273, 335
 - componential analysis of, 271
 - semantic differential, 270
 - and transfer, 301
 - in verbal learning, 282, 283, 286
- Mechanism, 357
- Mediation, language, 328
- Memorizing, 313
- Memory, 289, 307-335, 354, 359-360
 - consolidation, 322, 329
 - decay *vs.* interference, 318-320
 - echoic, 316
 - episodic, 308, 334-335
 - iconic, 313, 324, 335
 - interference, 308, 321-323, 331
 - long-term, 308, 313, 317, 320-323, 324, 335
 - primary, 316

Memory (*cont.*)

- running, 324
- semantic, 308, 334-335
- short-term, 308, 313, 316-320, 324, 335
- span of immediate, 282, 314, 316, 320, 327, 328, 345
- trace, 321, 330, 359
- working, 324
- Memory drum, 279
- Mentalism, 11-12, 357
- Mental representation, 346-347
 - enactive, iconic, and symbolic, 346
- Mental rotation, 340-341
- Metabolism, conditioning of, 206
- Metamemory, 324-325, 335
- Metaphor, 139, 240-241, 275, 347
 - search, 333, 343-345
 - storage and retrieval, 325-335
- Mill, James, 282
- Mind, 10, 12, 307, 357
- Mixed schedule, 187
- Mnemonic, 307, 309, 310-313, 325, 335, 339
 - method of loci, 311, 328, 331
 - peg-word technique, 312
- Molar theory, 112, 194
- Molecular theory, 112, 194
- Mongolian gerbil, 101
- Monkey, 49, 51, 53, 89, 101-102, 113, 163-165, 217, 304
- Monochromatic environment, 151
- Morse code, 235
- Motivation, 39, 60-61, 76, 226
- Motor memory, 334
- Motor skill, 119, 132, 283, 299, 346
 - discrete *vs.* continuous, 334
- Motor theory of thinking, 134
- Movement, 23-24, 65, 118, 339, 354
- Multiple schedule, 186-187, 195
 - and stimulus control, 140-141
- Music, 134-135
- Myna bird, 232fn
- Naming, 241
 - not tacting, 236
- Negation, 251
- Negative punishment, 113 (*see also* Punishment)
- Negative reinforcement, 104, 353 (*see also* Avoidance, Escape)
 - ambiguity of, 107-108
 - parameters, 105
 - vocabulary, 104
- Nervous system, 5, 40, 52
- Nonsense syllable, 243, 264, 270, 282
- Not-responding, as behavior, 114-115, 131
- Novelty, 137
 - and recognition, 296
 - in semantics, 290
 - in verbal behavior, 244, 252, 258, 274-275
- Observation, 18, 21, 24-26, 35, 226, 351
- Occasion (for response), 33, 36, 140, 165
- Oddity, 158
- Omission training, 113, 207-208, 213, 226, 353
- Operant, 32, 70, 273, 354 (*see also* Discriminated operant)
 - as class, 357-359
 - differentiation of, 119-123
 - and respondent, 199-200
- Operant level, 82
- Operation, 24-37, 222-223, 353-354 (*see also* Consequences, Observation, Presentation of stimulus, Stimulus control)
- Overlearning reversal effect, 303
- Overshadowing, 210-211
- Paired-associates learning, 243, 277, 280-281, 285-289, 297
 - and transfer, 299-301
- Paradigm, 63, 221-227
- Paramecium, 4-5
- Parameter, 105, 196, 217
- Paraphrase, 251, 263-264
- Parent, 23, 26-29, 58-59, 75, 94, 102, 115, 231, 241, 247
- Parrot, 232fn
- Pause:
 - post-reinforcement, 8, 182
 - in VR, 171
- Pavlov, I.P., 42, 47, 80, 120
- Pavlovian conditioning (*see* Conditioning)
- Peak shift, 149, 151
- Perception, 17
- Perceptual readiness or set, 339
- Performance, 358
 - group average *vs.* individual, 67
- Phoneme, 232, 243
- Phrase structure, 258-262
- Physiology, 5-6, 12fn, 40, 327, 356-357
- Pigeon, 8-10, 49-51, 102-103, 110, 118-119, 127-129
 - and autoshaping, 213-214
 - chamber, 68-70
 - memory, 310
 - and schedules, 171-197
 - and stimulus control, 144-154, 156-160
 - and tacting, 237-240, 252

- Place-learning, 146
- Plurality, 251, 260, 273, 274-275
- Poetry, 232
- Polydipsio, 49
- Population (*see* Class)
- Porpoise, 137
- Positive punishment (*see* Punishment)
- Positive reinforcement (*see* Reinforcement)
- Post-discrimination gradient, 149-151, 153
- Post-reinforcement pause, 8, 182
- Posture, 24, 40, 49, 118, 119, 207
- Potential, 39, 52-55
- Practice, 55-56
 - massed *vs.* spaced, 283
- Pragmatics, 255-256
- Preconditioning, sensory, 208-210, 226
- Predictability, in conditioning, 204-205, 214, 219
- Preference, 13, 20, 196
- Preparedness, 110, 160-162, 355 (*see also* Species-specific behavior)
- Presentation of stimulus, 24, 26-29, 35-37, 222-223 (*see also* Elicitation)
 - and stimulus control, 31-32, 199-215
- Primacy, 290, 291
- Primate, 25-26
- Prism, inverting, 89-90
- Private events:
 - as cognitive processes, 338, 341, 346
 - tacts of, 241-242
- Proaction, 277, 302-303
 - and memory, 308, 319, 323, 324, 333
- Probability, 42-45, 111, 117, 125, 199, 227
 - conditional, 44-45
 - in conditioning, 205, 219
 - of reinforcement, 172, 175-176
- Probe, 245
- Problem-box (Thorndike), 30, 64-65
- Problem-solving, 25-26, 29, 253, 330, 338, 346, 348-351, 360
- Process, 24-35, 75
- Processing, cognitive, 338-347
 - depth or level, 329
 - stages, 342-345
- Production, verbal, 245, 268
- Prompt, 243, 245
- Pseudoconditioning, 205fn
- Psychological reality, 261, 265-266
- Psychology:
 - behavioral, 11-12, 40, 42, 268, 338, 340, 357-360
 - cognitive, 11-12, 338-347, 357-360
 - gestalt, 330
- Punisher, 74, 223 (*see also* Punishment)
- Punishing stimulus (*see* Punisher)
- Punishment, 24, 37, 64fn, 75, 93-104, 120, 226, 337, 355, 360
 - and conditioning, 206
 - and discriminative stimuli, 102-103
 - effectiveness, 97
 - and elicitation, 100-102
 - negative, 113, 207-208
 - permanence, 97
 - and preparedness, 160
 - and reinforcement, 94-96
 - relativity, 99
 - schedules, 185, 192
 - vocabulary, 74fn, 94, 114
- Puzzle-box (Thorndike), 30, 64-65
- Quintilian, 311
- Raccoon, 136
- Rat, 18-21, 33-35, 48-49, 82-83, 119-124, 140-143, 215-221, 304
 - escape and avoidance, 57, 105-112
 - maze-learning, 65-70, 86-87, 133, 146
 - preparedness, 160-164
 - response hierarchy, 18-21, 60, 77-78, 99-100
- Rate, response, 68, 197
 - differentiation of, 131
- Ratio schedule (*see* FR, VR)
- Reaching, 33
- Reaction time, 340, 342
- Reading, 11-12, 156, 244, 247-250, 295, 301, 356, 357
 - not textual behavior, 234
- Recall, 117, 227, 309, 317, 360
 - cued, 277, 331
 - free, 280-281 (*see also* Free recall)
 - serial, 291
- Recency, 290, 291
- Recognition, 277, 280, 282, 297, 360
 - nonverbal, 295-296
 - and recall, 321
 - verbal, 295-297
- Recovery:
 - after punishment, 95-96
 - spontaneous, 80-82, 323
- Recursive structure, 266
- Redundancy, 283
- Reference, 229-230, 249-250, 255
- Reflex, 22, 39-48
 - alimentary, 31
 - Babinski, 43-44
 - conditional, 32, 40, 42, 45
 - laws of, 41
 - patellar, 40, 205
 - pupillary, 40, 206
 - salivary, 199, 205, 207

- Reflex (*cont.*)
 startle, 39-40, 51-53, 205fn
 walking, 57-59
- Reflex arc, 40
- Rehearsal:
 and iconic memory, 314-315
 and memory, 313, 324, 335, 344
 and short-term memory, 316-320
 in verbal learning, 291, 310
- Reinforcement, 8, 32, 37, 62-91, 93, 226, 350, 360
 and conditioning, 202
 continuous, 167, 184
 differential, 118-127 (*see also* Differential reinforcement)
 intermittent, 167
 negative, 93, 104-112
 parameters, 105
 partial, 167, 202
 permanence (*see* Extinction)
 and punishment, 94-95
 regular, 167
 relativity of (*see* Relativity)
 of verbal behavior, 244-245
 vocabulary, 73-76, 104
- Reinforcement schedules (*see* Schedules)
- Reinforcer, 74, 223, 226 (*see also* Reinforcement)
 conditioned, 61
- Reinforcing stimulus (*see* Reinforcer)
- Relation, discrimination of, 156-159, 239
- Relativity:
 in elicitation, 201-202
 of punishment, 99
 of reinforcement, 64, 76-78, 88, 136, 190, 225-226, 355
- Releaser, 28
- Remembering (*see* Memory)
- Reminiscence, 307, 322
- Repetition, 55-57
- Resistance to change, 79, 197
- Resistance to extinction, 79-80
- Respondent, 199-200, 358-359
- Respondent conditioning (*see* Conditioning)
- Response, 17, 22-23, 35, 39, 357
 aggressive, 82-83
 conditional (*see* CR)
 observing, 52, 157
 orienting, 52, 200-201
 unconditional (*see* UR)
- Response availability, 286-287
- Response hierarchy (*see* Hierarchy)
- Response integration, 286-287, 299, 303
- Response-learning, 65, 88, 146
- Response-shock (RS) interval, 108-109, 112
- Response strength, 131, 197
- Retention, 325, 329-330
- Retrieval, 308, 325-326, 329, 330-333
- Retroaction, 277, 302-303, 308, 323
- Reversal, 303-304
- Reversibility of reinforcement (*see* Relativity)
- Rhesus monkey, 217
- RS interval, 108-109, 112
- Running wheel, 19, 49, 77, 99
- Runway, 31, 67-68, 287
- Saccharin, 160-161
- St. Augustine, 330
- Salience, 143, 149, 210, 291
- Salivation, 31, 40-41, 57, 121, 199, 200, 205
 operant, 207-208
 spontaneous, 121, 200
- Satiation, 118, 199
 semantic, 271
- Savings, method of, 285, 320, 321
- Scallop, 72, 180-181, 186, 247
 inverted, 185
- Schedules, 129-131, 167-197
 classified, 183-187
 and gradients, 147
- Search:
 exhaustive *vs.* self-terminating, 343-344
 memory, 333, 342-345
 random *vs.* fixed-order, 343, 344
 serial *vs.* parallel, 343-345
 visual, 292, 339
- Second-order schedule, 183
- Selection (*see* Evolution)
- Self-control, 196
- Self-embedded structure, 267
- Self-injury, 102, 115
- Semantic differential, 270
- Semantics, 250, 255, 263, 268-275, 337, 359-360
 and memory, 308, 319-320, 327
 and syntax, 269
 and verbal learning, 292-294
- Sensitization, 53fn
- sensory-motor learning, 64, 88-91, 226, 353
- Sensory preconditioning, 208-210, 226
- Sensory processes, 88-91, 338-342
- Sensory reinforcement, 89
- Sentence, 255, 270
 ambiguous, 258, 260, 268
- Sequence, verbal (*see* Intraverbal, Serial learning)
- Sequential behavior, 10, 18, 42, 85-86, 132
 and acquisition, 164-165
 chains *vs.* temporally extended units, 133-135
- Sequential dependencies in language, 256-258, 267
- Serial anticipation, 280, 283

- Serial learning, 243, 277, 279-280, 282-286, 288, 297
- Serial-position effect, 283-284, 289
- Serial recall, 279-283
- Shaping, 34, 117-119, 125-127, 213, 360
 - and fading, 154
 - impossibility in conditioning, 200
 - and verbal behavior, 232, 246
 - of visualizing, 341
 - vocabulary, 155
- Shift:
 - extradimensional *vs.* intradimensional, 304
 - peak, 149, 151
 - reversal *vs.* nonreversal, 304
- Shock (*see* Electric shock)
- Shock-shock (SS) interval, 108-109, 112
- Side effects:
 - of extinction, 83
 - of punishment, 99-103
- Signal, 24, 31-32, 36-37, 199, 354
 - warning, 57, 111, 220
- Signalled avoidance, 57
 - two-process theory, 110-112
- Sign language, 231, 248
- Sign-tracking, 214
- Similarity, 282, 283, 300-301, 305
- Simile, 240
- Simonides, 311
- Skinner, B. F., 241
- Sleep, 25, 283, 321, 329
- Slope of cumulative record, 70-71
- Spaced practice, 283
- Span of immediate memory (*see* Memory)
- Speaking, 229, 231-234, 246
- Species, 358-360
- Species-specific behavior, 13, 26-28, 70, 101, 110, 136, 206, 354-359
- Species-specific defense reactions, 110, 206
- Spirit, 353
- Spontaneous behavior, 56-57, 121
- Spontaneous recovery, 80-82
- S-R bonds, 134, 135
- SS interval, 108-109, 112
- State-dependency, 308, 333
- Steady-state, 164, 196
- Stereotypy of responses, 84
- Stimulus, 17, 22-23, 35, 39, 78, 156, 222, 357
 - appetitive, 20, 223, 226
 - aversive, 20, 61, 93, 114, 206, 223, 226
 - comparison, 157-158
 - conditional (*see* CS)
 - contingent, 221-224
 - discriminative (*see* Discriminative stimulus)
 - functional, 278, 283
 - neutral, 20, 88, 201-203, 224, 226
 - nominal, 278
 - preaversive, 216-221
 - sample, 156-158
 - supernormal, 27
 - unconditional (*see* US)
 - warning, 57, 111, 220
- Stimulus control, 24, 222-223, 337, 351, 354, 356 (*see also* Discriminated operant, Discrimination, Gradient, Multiple schedule)
 - and attention, 143-145, 153, 154
 - in conditioning, 199-215
 - and consequences, 32-37
 - and elicitation, 31-32, 35-37
 - and fading, 153-155
 - in tacting, 236-242, 256
 - and verbal learning, 286, 289
- Stimulus-learning, 65, 88
- Stimulus-presentation operation (*see* Presentation of stimulus)
- Stimulus-response bonds, 134, 135
- Storage, 308, 325-329
- Straight alley, 31, 67-68, 287
- Strain, ratio, 10, 173
- Strength, reflex, 41-42
- Strength, response, 131, 197
- Structure, 11-12, 23, 28, 137, 183, 337, 354, 356-358
 - deep *vs.* surface, 263
 - hierarchical, 266-267, 271-272, 346-347
 - of lists, 289
 - and memory, 308, 327, 335
 - phrase, 258-262
 - and problem-solving, 348
 - of verbal behavior, 230, 243, 255-273
- Subvocal responses, 233, 234
- Successive approximations (*see* Fading, Shaping)
- Sucking, 55-56
- Superstition, 50-51, 177-180
 - and extinction, 83-85
- Suppression, conditioned, 217-221
- Symbolic matching, 158
- Syntax, 255-268, 274-275, 337, 359-360
 - and memory, 308, 334
 - and semantics, 269
- Tact, 236-243, 246, 248-249, 360
 - abstract, 237-240, 347
 - constructed, 253
 - extended, 240-241
 - and memory, 309-310, 325
 - of relations, 239, 251
 - of verbal dimensions, 256, 260
- Tandem schedule, 187, 190, 218
- Temperature, 5, 108, 206, 354
- Tense, 240, 252

- Textual behavior, 233-234, 236, 240, 242, 301
- Theophrastus, 326
- Theory, 15, 22, 359
- language, 230
 - memory, 313, 321-323
 - molar vs. molecular, 112, 194
 - two-process, of avoidance, 110-112
- Thinking, 234, 253, 338, 354
- Thorndike, E. L., 96-97
- Threshold, 41, 110, 150, 239
- Thumb twitch, 127-128
- Time, discrimination of, 181
- Time, tact of, 240
- Time-out from positive reinforcement, 113
- Tip-of-the-tongue phenomenon, 11, 239, 325
- Topographical drift, 51, 84
- Topography, 51, 119, 127, 214, 255, 338, 359
- Toxicosis, 162
- Transcription, 232-233, 236, 301
- Transfer, 139, 158, 164, 248, 277, 297-305
- and memory, 308, 323
 - negative, 298-305, 350
 - nonspecific, 299
 - nonverbal, 303-304
 - positive, 298-305
 - and problem-solving, 348-350
 - zero, 299
- Transfer surface, 300-301
- Transformation, 256, 263-267
- obligatory vs. optional, 266
- Translation, 235
- Transportation, 304
- Trial-and-error, 64
- Truth, 7, 242, 246
- Two-process theory of avoidance, 110-112
- Typing, 233
- Unconditional response (*see* UR)
- Unconditional stimulus (*see* US)
- Understanding, 250
- reading for, 234
- Unfinished tasks, 325
- Unit, behavioral, 132 (*see also* Discriminated operant, Operant, Respondent)
- in memory, 337
 - in problem-solving, 348
 - reflex, 42, 47
 - in schedules, 130, 182-183
 - temporally extended, 133-135
 - verbal, 232-235, 243, 252, 268, 275, 288, 297
- UR, 201-203, 208-215
- US, 201-203, 208-215
- Variability, 84, 358
- and causation, 177
 - differentiation of, 137
 - in encoding, 329
- Variable-interval schedule (*see* VI schedule)
- Variable-ratio schedule (*see* VR schedule)
- Variable-time schedule (*see* VT schedule)
- Variation, 358
- Verb, 260
- Verbal behavior, 10, 134-135, 137, 229-253, 354, 360
- Verbal community, 239, 241
- Verbal discrimination, 277, 280-281, 292-295, 297
- conditional, 293-294
- Verbal learning, 96, 277-297 (*see* Free recall, Paired-associates learning, Recognition, Serial learning, Verbal discrimination)
- VI schedule, 167-180, 183, 184, 216, 218
- extinction after, 172-173
 - and gradients, 147, 152
 - and IRTs, 175-177
 - in multiple schedule, 140-141
 - yoked to VR, 174-175
- Vision, 229, 327, 337, 340-342, 353
- inverted, 89-90
- Visual coordination, 29, 33, 89-91
- Vitalism, 357
- Vocal behavior, 231-232, 274, 279
- VR schedule, 167-180, 184
- extinction after, 172-173
 - and IRTs, 175-177
 - yoked to VI, 174-175
- VT schedule, 178, 183, 184
- Walking, 57-59
- Water-jar problem, 347-350
- Watson, J. B., 42, 47, 340
- Wittgenstein, L., 227, 241
- Word, 229, 249, 270
- Writing, 229, 232-236, 279
- X-irradiation, 160-161
- Yoked schedules, 174-175, 247
- Zebra danio, 55



Form No. 3.

PSY, RES.L-1

**Bureau of Educational & Psychological
Research Library.**

The book is to be returned within
the date stamped last.

.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		
.....		

WBGP-59/60-5119C-5M